Power and limitations of convex formulations via linear and semidefinite programming lifts

by

Hamza Fawzi

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2016

Author .....
Department of Electrical Engineering and Computer Science
August 3, 2016

Certified by .....
Pablo A. Parrilo
Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by .....
Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chairman, Department Committee on Graduate Theses

Power and limitations of convex formulations via linear and semidefinite programming lifts

by

Hamza Fawzi

Submitted to the Department of Electrical Engineering and Computer Science
on August 3, 2016, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Convex relaxation methods play an important role in mathematical optimization to tackle hard nonconvex problems, and have been applied successfully in many areas of science and engineering. At the heart of such methods lies the question of obtaining a tractable description of the convex hull of a set. In this thesis we focus on the question of finding tractable representations of convex sets via the method of lifting, whereby the "hard" convex set is expressed as the projection of a simpler one living in higher-dimensional space. We derive new results and insights on the power and limitations of such liftings.

In the first part of the thesis we study limitations of the lifting method and develop lower bounds on the sizes of linear programming (LP) and semidefinite programming (SDP) lifts of polytopes. For LP lifts the bound we develop applies generally for the nonnegative rank of matrices and we compare our method with existing combinatorial and non-combinatorial techniques. For SDP lifts we focus on so-called *equivariant lifts* that respect symmetry, and obtain lower bounds on the size of such lifts for certain combinatorial polytopes by exploiting the connection with the sum-of-squares method.

In the second part of the thesis, we study the power of the lifting procedure and show how to obtain small semidefinite lifts for certain classes of polytopes via the idea of sparse sums of squares. We develop a graph-theoretic method to construct such lifts and use it to resolve a conjecture of Laurent from 2003 on the cut polytope, and to give an explicit sequence of polytopes with a gap between LP and SDP lifts.

Finally we depart from the specific question of constructing lifts and consider the general problem of certifying nonnegativity of functions. We study a class of certificates rooted in convex duality and show that they encompass many existing methods for proving nonnegativity based on convex optimization. In particular we propose a new proof system to certify nonnegativity of entropy-like functions, which we illustrate on the problem of computing the logarithmic Sobolev constant of finite Markov chains.

Thesis Supervisor: Pablo A. Parrilo
Title: Professor of Electrical Engineering and Computer Science

Acknowledgments

I was very fortunate to be advised by Pablo Parrilo during my PhD and it is my great pleasure to thank him for his continued support and encouragements. Throughout our discussions and meetings his insights and intuition have shaped my understanding on many different topics in convex optimization and applied mathematics in general. I want to thank him for the energy and passion he always brings to our meetings, and for giving me the freedom in exploring different research directions during my PhD. I would also like to thank Profs. Ankur Moitra and John Tsitsiklis for being on my thesis committee.

My thanks also go to Paulo Tabuada, my Master's thesis advisor at UCLA, who gave me the opportunity to visit MIT during my time at UCLA for his support and encouragements.

Many of the results in this thesis have been obtained in collaboration with James Saunderson. It is a pleasure to thank James for all the work that we have done together and for everything I learned from him. I have really enjoyed all the meetings that we had and I look forward to future collaboration. I would also like to thank my officemates, Omer, Noele, Quan, and Jennifer, and friends at LIDS as well as all the staff for the great working environment that I enjoyed a lot. My thanks also go to all my friends from the Muslim Students Association for the great times and enjoyable moments we spent together.

I would also like to thank my brothers Omar and Hussein for the various discussions we had on different topics ranging from quantum information theory to computer vision and deep learning that I have really enjoyed and that were very beneficial. Finally I would like to thank my parents for their continued support and encouragements.

The research in this thesis was funded in part by grants AFOSR FA9550-11-1-0305 and AFOSR FA9550-12-1-0287.

Contents

Chapter 1

Introduction

In this short introductory chapter we give a brief overview of the lifting method and discuss its importance in optimization when combined with the idea of *convex reformulations*. We illustrate the power of the lifting method on a simple example and introduce in an informal way linear programming and semidefinite programming lifts. We briefly discuss the history of the lifting method in optimization, and conclude by presenting the organization of the thesis.

1.1 Convex reformulations

Consider the following problem where we want to minimize a linear function $\ell(x)$ subject to the constraint $x \in X$:

$$\text{minimize} \quad \ell(x) \quad \text{subject to} \quad x \in X. \tag{1.1}$$

Here X is an arbitrary subset of $\mathbb{R}^n$ and need not be convex. It is a well-known fact that, since the objective function is linear, the optimal value of (1.1) remains unchanged if we change the constraint "$x \in X$" by "$x \in \text{conv}(X)$", where $\text{conv}(X)$ denotes the convex hull of X (see Figure 1-1 for an illustration):

$$\text{minimize} \quad \ell(x) \quad \text{subject to} \quad x \in \text{conv}(X). \tag{1.2}$$

Recall that the convex hull of X is the set of all possible convex combinations of elements of X:

$$\text{conv}(X) = \left\{ \sum_{i=1}^{m} \lambda_i x_i : m \in \mathbb{N}, \lambda_1, \ldots, \lambda_m \geq 0, \sum_{i=1}^{m} \lambda_i = 1, x_1, \ldots, x_m \in X \right\}.$$

Using this definition the equality of the optimal values of (1.1) and (1.2) is straightforward to verify. Note that problem (1.2) is now formally *convex* since the cost function is linear and the feasible set is convex.

The transformation from (1.1) to (1.2) seems to rely heavily on the fact that ℓ is a linear function. It turns out however that a similar transformation can be applied

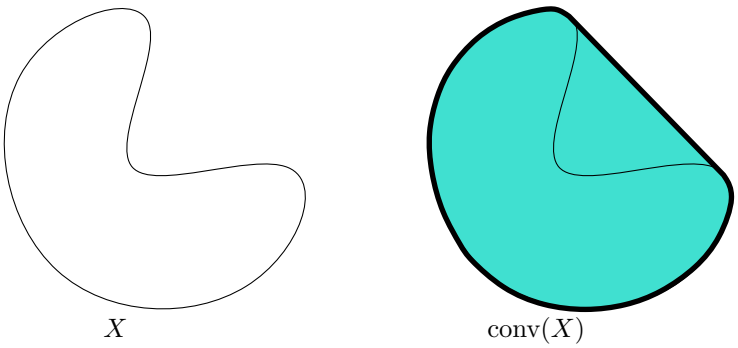

Figure 1-1: A nonconvex set X and its convex hull

more generally, even if the objective is not linear, by introducing additional variables and constraints. To illustrate this, assume that our objective in (1.1) was quadratic instead of being linear, i.e., we are interested in minimizing $q(x) = \sum_{i \leq j} q_{ij} x_i x_j$ subject to $x \in X$:

$$\text{minimize} \quad q(x) = \sum_{i \leq j} q_{ij} x_i x_j \quad \text{subject to} \quad x \in X. \tag{1.3}$$

If we introduce additional variables y_{ij} playing the role of $x_i x_j$ we can reformulate the problem above as follows:

$$\text{minimize} \quad \sum_{i \leq j} q_{ij} y_{ij} \quad \text{subject to} \quad y \in Y \tag{1.4}$$

where Y is defined as

$$Y = \{(y_{ij})_{i \leq j} : \exists x \in X \text{ s.t. } y_{ij} = x_i x_j \; \forall 1 \leq i \leq j \leq n\}.$$

The objective function of (1.4) is now linear and thus by the same reasoning as above the constraint "$y \in Y$" can be changed to "$y \in \text{conv}(Y)$".

In both cases we have transformed the original problem to a new problem that is convex, at least formally. In order to solve the problem however we need to find a tractable representation of the set $\text{conv}(X)$ (or $\text{conv}(Y)$).

Note that the idea of *convex reformulations* has been recognized in mathematical optimization since the early days of integer programming. We refer the reader to [57] and the references therein for more details on the use of such reformulations in combinatorial optimization.

1.2 Lifts of convex sets

In this thesis we will be mostly dealing with the case where the set X is finite which arises in discrete and combinatorial optimization. The corresponding convex set $P = \text{conv}(X)$ in this case is called a *polytope* and can be described using a finite number

of linear inequalities.

The problem of optimizing a linear function over a polytope P is known as linear programming. Interior-point methods are a popular class of algorithms for linear programming, and the complexity of such algorithms typically depend on the size of the inequality description of the polytope[1]. The size of the trivial such description is equal to the number of *facets* of P, a geometric quantity associated to P. Unfortunately, in many cases of interest, the number of facets of P is prohibitively large to enumerate directly in a linear programming formulation.

The idea of *lifting* consists in expressing the polytope P as the projection of a higher-dimensional polytope Q that has much fewer facets than P. We say in this case that Q is a (linear programming) *lift* of P. For the purpose of optimization one can then work over Q rather than working over P. Indeed if we are interested in minimizing ℓ over P and if $P = \pi(Q)$ where π is a linear (projection) map then we have:

$$\min_{x \in P} \ell(x) = \min_{y \in Q} \ell \circ \pi(y). \tag{1.5}$$

To give a simple example of a lift let P denote the ℓ_1 ball in $\mathbb{R}^n$:

$$P = \{x \in \mathbb{R}^n : \|x\|_1 \leq 1\}.$$

The trivial description of P obtained by enumerating the facets has size 2^n since the facets of the ℓ_1 ball are given by all possible inequalities of the form $\pm x_1 \pm x_2 \pm \cdots \pm x_n \leq 1$. It is not difficult however to see that the ℓ_1 ball admits a description using only $2n$ linear inequalities, namely:

$$P = \left\{ x \in \mathbb{R}^n : \exists y \in \mathbb{R}^n \text{ s.t. } -y_i \leq x_i \leq y_i, \sum_{i=1}^{n} y_i = 1 \right\}. \tag{1.6}$$

In this case the higher-dimensional polytope Q lives in $\mathbb{R}^{2n}$ and consists of vectors (x, y) that satisfy the constraints $-y_i \leq x_i \leq y_i$ and $\sum_{i=1}^{n} y_i = 1$. The projection map is $\pi : (x, y) \mapsto x$. The key point to note here is that the number of inequalities in the description (1.6) is now $2n$, instead of the 2^n we had before. Observe that the description (1.6) is the one that we often use in practice when solving ℓ_1 optimization problems. In Chapter 2 we give other examples of polytopes where lifting allows us to get a description that is much smaller than the trivial one.

So far we have been interested in expressing the polytope P as the projection of another polytope Q. A significant portion of this thesis will deal with a more general class of lifts called *semidefinite programming* lifts. Here the goal is to express the polytope P as the projection of the feasible set of a *semidefinite program* (SDP), i.e., a convex set Q that can be described using linear matrix inequalities:

$$Q = \{y \in \mathbb{R}^m : F_0 + y_1 F_1 + \cdots + y_m F_m \succeq 0\}. \tag{1.7}$$

[1]For a more detailed discussion of complexity-theoretic aspects we refer the reader to Chapter 2, Section 2.4.

$F_0, \ldots, F_m$ are $d \times d$ real symmetric matrices and the constraint in (1.7) indicates that the matrix $F_0 + y_1 F_1 + \cdots + y_m F_m$ is positive semidefinite. Such a convex set is sometimes called a *spectrahedron*. When the matrices $F_0, \ldots, F_m$ are diagonal the set Q is a polytope; in general however Q is not necessarily a polytope. It is clear from this observation that SDP lifts form a broader class than LP lifts. The question of when one can find SDP lifts that are significantly smaller than LP lifts is still not very well understood. One of the results proved in Chapter 5 of this thesis shows that there is an explicit class of polytopes for which SDP lifts are vanishingly smaller than any LP lift (for increasing dimensions).

1.3 History

The idea of lifts (also called *extended formulations*) which consists in lifting the problem to a higher-dimensional space by introducing additional variables is well known in optimization. However the first paper that studies lifts in a systematic way to prove nonexistence of small lifts for certain polytopes is due to Yannakakis in 1991 [101]. In his paper Yannakakis showed that the traveling salesman polytope and the matching polytope do not have polynomial-size symmetric linear programming lifts.

The recent years have witnessed a resurgence of interest in this topic. From the lower bounds point of view several results have been proved concerning the nonexistence of small lifts for polytopes arising in combinatorial optimization. Fiorini et al. [41] resolved a conjecture left open by Yannakakis and proved that the traveling salesman polytope does not admit any polynomial-size linear programming lifts (without any symmetry requirement). Later, Lee, Raghavendra, Steurer [76] showed that the traveling salesman polytope has no polynomial-size semidefinite programming lift. Another major result was also obtained recently by Rothvoß [89] where he showed that the matching polytope has no polynomial-size linear programming lift (again, with no symmetry requirement). This result of Rothvoß is particularly striking since the matching polytope is known to have a polynomial-time separation oracle. Several results have also been obtained concerning approximate lifts, see for example [15, 21, 18].

From the upper bounds point of view new methods have been proposed to construct improved lifts for certain classes of polytopes of convex sets, see for example [62, 47, 64, 38, 94]. The problem of constructing semidefinite programming lifts for algebraic sets has been of specific importance in the area of *convex algebraic geometry* [10, 52, 69]. A conjecture by Helton and Nie [59] states that any convex semialgebraic set admits a semidefinite programming lift.

1.4 Organization

The thesis is organized as follows:

- Chapter 2 starts by giving the formal definitions of LP and SDP lifts. We then present a systematic way to understand lifts of polytopes in terms of *certifi-*

cates of nonnegativity of facet inequalities. Such a characterization is due to Yannakakis [101] in the case of LP lifts and to Gouveia, Parrilo, Thomas [50] for the case of SDP lifts (and more generally conic lifts). This point of view on lifts is crucial for the rest of the thesis.

- Chapter 3 considers the problem of lower bounding the *nonnegative rank* of a matrix. As we see in Chapter 2 the nonnegative rank plays an important role in characterizing the size of the smallest LP lift of a polytope. Several techniques have been proposed in the literature to lower bound the nonnegative rank. In this chapter we first review the different techniques and we present a new method that unifies some of the existing techniques. This chapter is based on the paper [34].

- Chapter 4 is devoted to the study of so-called *equivariant SDP lifts*, which are lifts that respect the symmetries of the original polytope. We derive a structure theorem that gives a characterization of such lifts in terms of sum-of-squares certificates of facet inequalities from an invariant subspace. We apply our structure theorem to derive lower bounds for the cut polytope, the parity polytope, and regular polygons in the plane. This chapter is based on the papers [37, 36].

- Chapter 5 is concerned with constructing semidefinite programming lifts by exploiting the idea of *sparse* sums of squares. By working in a general setting of Fourier analysis on finite abelian groups and by exploiting certain results on sparse positive semidefinite matrices we show that there exists a family of polytopes in increasing dimensions with a growing gap between LP and SDP lifts. The tools we develop also allow us to prove a conjecture of Laurent from 2003 [72] on the Lasserre hierarchy for the maximum cut problem. This chapter is mostly based on the paper [38].

- Finally in chapter 6 we depart from the specific problem of constructing lifts of polytopes and we consider more generally the problem of certifying nonnegativity of a function on a given set. We first show how ideas from [50] allow us to formulate certificates of nonnegativity that generalize the existing ones (LP, SDP/SOS, geometric programming, signomials, etc.). We then use this framework to develop new certificates of nonnegativity for a class of entropy-like functions that cannot be handled using existing techniques. As an application of our method we show how it can be used to obtain a numerical estimate of the logarithmic Sobolev constant for any given finite Markov chain.

The results in Chapter 2 concerning the characterization of lifts are used in Chapters 4 and 5. Chapter 3 can be read independently and only uses the definition of nonnegative rank (which we recall anyway at the beginning of the chapter). Chapter 6 is independent of the other chapters though many of the ideas presented there are inspired from results in Chapters 2, 4 and 5.

1.5 Terminology and notations

The following table summarizes some of the common notations used throughout the thesis. More specific notations will be defined in the individual chapters.

$\mathbb{R}_+$ (resp. $\mathbb{R}_{++}$)	nonnegative (resp. positive) real numbers
$\mathbf{S}^d$	space of $d \times d$ real symmetric matrices
$\mathbf{S}^d_+$ (resp. $\mathbf{S}^d_{++}$)	cone of $d \times d$ real symmetric positive semidefinite (resp. positive definite) matrices
X^* (for $X \in \mathbb{C}^{n\times m}$)	Hermitian conjugate of X defined by $(X^*)_{ij} = \overline{X_{ji}}$
$\mathbf{H}^d$	space of $d \times d$ complex Hermitian matrices
$\mathbf{H}^d_+$ (resp. $\mathbf{H}^d_{++}$)	cone of $d \times d$ Hermitian positive semidefinite (resp. positive definite) matrices
$\mathbf{S}^V, \mathbf{S}^V_+, \mathbf{H}^V, \mathbf{H}^V_+$	same as above except that rows and columns are indexed by some set V
$\mathbf{S}^d_+$	cone of $d \times d$ real symmetric positive semidefinite matrices
E^* (for finite dim. vector space E)	space of linear forms on E = dual space of E
K^* (for cone $K \subseteq E$)	$\{\ell \in E^* : \ell(x) \geq 0 \; \forall x \in K\}$ = dual cone of K (if E has inner product $\langle \cdot, \cdot \rangle$ can identify K^* as a cone in E)
$\mathbb{R}[x_1, \ldots, x_n]$	space of polynomials in n variables $x_1, \ldots, x_n$
$\mathbb{R}[x_1, \ldots, x_n]_{\leq k}$	space of polynomials of degree at most k in n variables $x_1, \ldots, x_n$

Chapter 2

Lifts of convex sets and certificates of nonnegativity

The main goal of this chapter is to give a concrete way to think about lifts of a polytope in terms of *certificates of nonnegativity of its valid linear inequalities*. Any polytope P (or more generally, any closed convex set) is characterized by its set of valid linear inequalities: these are the affine functions that take nonnegative values on P. A key result of Yannakakis [101], extended later in [50], shows that producing a lift of P is equivalent to finding *certificates of nonnegativity* of all the valid linear inequalities of P. What distinguishes LP lifts from SDP lifts is the kind of certificates of nonnegativity considered. This point of view on lifts will be crucial for the rest of the thesis and this chapter is thus devoted to explaining and illustrating it. We will also outline connections with certain matrix factorization problems.

2.1 LP lifts

A polyhedron in $\mathbb{R}^N$ is a set described using a finite number linear inequalities:

$$Q = \{x \in \mathbb{R}^N : b + Ax \geq 0\} \tag{2.1}$$

where $A \in \mathbb{R}^{d \times N}$ and $b \in \mathbb{R}^d$. Equivalently a polyhedron can be described in standard form as the intersection of the nonnegative orthant $\mathbb{R}^d_+$ with an affine subspace L.

Definition 1. Let P be a polytope in $\mathbb{R}^n$. We say that P has a *LP lift of size d* if P can be written as $P = \pi(\mathbb{R}^d_+ \cap L)$ where $\pi : \mathbb{R}^d \to \mathbb{R}^n$ is a linear map and L is an affine subspace of $\mathbb{R}^d$. The size of the smallest LP lift of P is called the *LP extension complexity* of P and denoted $\mathrm{xc}_{\mathrm{LP}}(P)$.

Remark 1. Equivalently, an LP lift of size d for a polytope P is a representation $P = \pi(Q)$ where Q is a polytope with d facets.

Figure 2-1 illustrates an LP lift of size 5 for the regular hexagon in the plane.

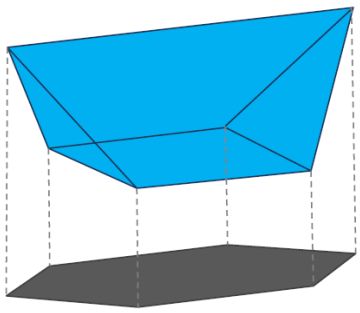

Figure 2-1: LP Lift of a hexagon of size 5. Note that hexagon has 6 facets whereas the higher-dimensional polytope has 5 facets in $\mathbb{R}^3$.

2.1.1 Examples of LP lifts

We now give some examples of polytopes P that admit nontrivial lifts.

- We saw in Chapter 1 the example of the ℓ_1 ball in $\mathbb{R}^n$ which has 2^n facets and which admits a simple lift of size $2n$.
- Another example of nontrivial lift is for the *permutahedron*. The permutahedron $P \subset \mathbb{R}^n$ is defined as the convex hull of all possible permutations of the vector $(1, 2, \ldots, n)$, i.e.:

  $$P = \mathrm{conv}\,\{(\sigma(1), \ldots, \sigma(n)) : \sigma \in \mathfrak{S}_n\}$$

  where $\mathfrak{S}_n$ is the set of permutations on $\{1, \ldots, n\}$. This polytope arises naturally in ordering problems such as in gene sequencing, see e.g. [79]. It is known that the permutahedron has an exponential number of facets, precisely $2^n - 2$.

However it is not very difficult to construct a lift of the permutahedron of size n^2. Indeed let Q denote the convex hull of *permutation matrices* in $\mathbb{R}^{n\times n}$. The Birkhoff-von Neumann theorem asserts that Q is precisely the set of doubly stochastic matrices, i.e.:

$$Q = \Big\{ M \in \mathbb{R}^{n\times n} \ : M_{ij} \geq 0 \ \ \forall i,j = 1,\dots,n, \\ \sum_{i=1}^{n} M_{ij} = 1, \ \forall j = 1,\dots,n, \\ \sum_{j=1}^{n} M_{ij} = 1, \ \forall i = 1,\dots,n \Big\}.$$

It is easy to see that P is a projection of Q: indeed if we let $\pi : \mathbb{R}^{n\times n} \to \mathbb{R}^n$ defined by $\pi(M) = Mu$ where $u = (1,2,\dots,n)^T$ then we get that $\pi(Q) = P$. This lift has size n^2 because Q requires exactly n^2 inequalities for its description. It turns out however that this lift is not optimal. Goemans showed in [47] that the permutahedron admits a lift of size $O(n \log n)$ which is optimal (i.e., there is no smaller possible lift). His construction however is more complicated and makes use of sorting networks.

- Let $P \subset \mathbb{R}^2$ be the regular N-gon in the plane, i.e., the convex hull of the N complex roots of unity. Even though P has N facets (and N vertices), a result of Ben-Tal and Nemirovski [6] shows that the regular N-gon admits a lifted description with only $O(\log N)$ inequalities. This construction was used in [6] to obtain polyhedral approximations of the second-order cone. Figure 2-1 shows a lift of the regular hexagon of size 5.

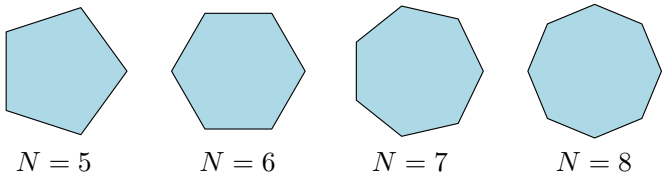

Figure 2-2: Ben-Tal and Nemirovski showed in [6] that the regular N-gon admits a LP lift of size $O(\log N)$.

- For other examples of lifts from the combinatorial optimization literature we refer the reader to the surveys [27, 61].

2.1.2 Yannakakis' theorem

Any polytope $P \subset \mathbb{R}^n$ (in fact any closed convex set) is described by its set of valid linear inequalities: these are the affine functions on $\mathbb{R}^n$ that are nonnegative on P.

Polytopes are described by a finite number of such linear inequalities known as the *facet inequalities*. Any facet inequality takes the form

$$\ell(x) \leq \ell_{\max}$$

where ℓ is a linear function and $\ell_{\max} := \max_{x \in P} \ell(x)$. For example in Figure 2-3 we show a facet inequality of the regular hexagon: in this case $\ell(x) = x + \frac{1}{\sqrt{3}} y$ and $\ell_{\max} = 1$.

Let X denote the vertices of the polytope P and let $\ell \leq \ell_{\max}$ be a facet inequality for P. Note that $\ell_{\max} - \ell|_X$ (i.e., the restriction of $\ell_{\max} - \ell$ to X) is a nonnegative function on X. The next theorem, due to Yannakakis shows that producing a lift of P is equivalent to finding *certificates of nonnegativity* for $\ell_{\max} - \ell|_X$.

Theorem 1 (Yannakakis [101])**.** *Let $P = \mathrm{conv}(X) \subset \mathbb{R}^n$ be a full-dimensional*[1] *polytope. Then P has an LP lift of size d if and only if, there exist d nonnegative functions on X, $a_1, \ldots, a_d : X \to \mathbb{R}_+$ such that the following holds: for any facet inequality $\ell(x) \leq \ell_{\max}$ of P there exist nonnegative coefficients $b_1, \ldots, b_d \geq 0$ such that*

$$\ell_{\max} - \ell|_X = \sum_{i=1}^{d} b_i a_i. \tag{2.2}$$

Note that Equation (2.2) is an equality of functions on X. It should be interpreted as a certificate of nonnegativity of $\ell_{\max} - \ell$ on X: indeed the right hand side is a linear combination with nonnegative weights (the b_i's) of nonnegative functions $a_1, \ldots, a_d$, and so is "obviously" nonnegative on X. The size of the certificate, here the number of functions $a_1, \ldots, a_d$, gives us the size of the LP lift. The main question in constructing a lift is of course to come up with the nonnegative functions $a_1, \ldots, a_d$.

Remark 2. In the statements above, a "*function on X*" is simply an element of $\mathbb{R}^X$, i.e., it can be simply seen as a vector of length $|X|$. We use the terminology "*function on X*" because in later chapters the domain X and its symmetries will play an important role when understanding so-called *equivariant* lifts.

Example 1. To illustrate Theorem 1 we now give the functions $a_1, \ldots, a_d$ that correspond to the LP lift of size 5 of the hexagon from Figure 2-1. For the regular hexagon we have $|X| = 6$ and so we will represent functions on X as column vectors of size 6 where the vertices are ordered counter-clockwise starting from the point with coordinates $(1, 0)$. The matrix shown in Figure 2-3 (left) gives the 6 facets of the regular hexagon, one facet per column. For example to see how the first column is formed note that the equation of the first facet inequality of the hexagon is $1 - x - y/\sqrt{3} \geq 0$. The k'th vertex of the hexagon has coordinates $(\cos(2(k-1)\pi/6), \sin(2(k-1)\pi/6))$. The k'th component of the first column is thus given by $1 - \cos(2(k-1)\pi/6) - \sin(2(k-1)\pi/6)/\sqrt{3}$. Note that each column has exactly two zeros: this is because each facet of the regular hexagon passes through exactly two vertices.

[1] We assume throughout this thesis that the polytope P of interest is full-dimensional. Some of the results may hold in greater generality but we keep this assumption for convenience.

$$\left[\begin{array}{c|c|c|c|c|c} 0 & 1 & 2 & 2 & 1 & 0 \\ 0 & 0 & 1 & 2 & 2 & 1 \\ 1 & 0 & 0 & 1 & 2 & 2 \\ 2 & 1 & 0 & 0 & 1 & 2 \\ 2 & 2 & 1 & 0 & 0 & 1 \\ 1 & 2 & 2 & 1 & 0 & 0 \end{array}\right] \qquad (2.3)$$

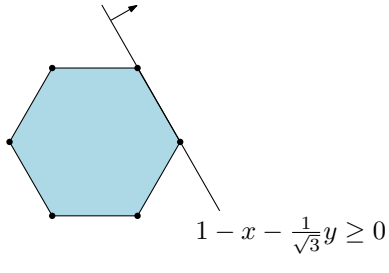

Figure 2-3: Facets of the regular hexagon

Consider now the following nonnegative functions on X, one per column (we do not explain now how we come up with such functions, in fact this is the difficulty in constructing lifts):

$$\left[\begin{array}{c|c|c|c|c} a_1 & a_2 & a_3 & a_4 & a_5 \\ \hline 0 & 1 & 0 & 0 & 1 \\ 0 & 2 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 2 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 \end{array}\right]. \qquad (2.4)$$

We need to show that these functions $a_1, \ldots, a_5$ satisfy the condition of Theorem 1, namely that each column of the matrix (2.3) is a nonnegative combination of the columns of (2.4). One can verify that this is true and that the coefficients (the b_i's in the notations of Theorem 1) are given in the right-most matrix below:

$$\underbrace{\left[\begin{array}{c|c|c|c|c|c} 0 & 1 & 2 & 2 & 1 & 0 \\ 0 & 0 & 1 & 2 & 2 & 1 \\ 1 & 0 & 0 & 1 & 2 & 2 \\ 2 & 1 & 0 & 0 & 1 & 2 \\ 2 & 2 & 1 & 0 & 0 & 1 \\ 1 & 2 & 2 & 1 & 0 & 0 \end{array}\right]}_{\text{the facets}} = \underbrace{\left[\begin{array}{c|c|c|c|c} 0 & 1 & 0 & 0 & 1 \\ 0 & 2 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 2 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 \end{array}\right]}_{\text{the functions } a_1, \ldots, a_5} \underbrace{\left[\begin{array}{c|c|c|c|c|c} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 & 2 \\ 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 2 & 1 & 0 & 0 \end{array}\right]}_{\text{the coefficients } b_i\text{'s}} \qquad (2.5)$$

$\Diamond$

2.1.3 Nonnegative matrix factorization

It is not difficult to see from the previous example that Yannakakis' theorem can be expressed in terms of nonnegative factorization of matrices. We recall now the definition of a *nonnegative factorization* of a matrix.

Definition 2 (Nonnegative matrix factorization / nonnegative rank)**.** Let $S \in \mathbb{R}_+^{p \times q}$ be a matrix with nonnegative entries. We say that S has a *nonnegative factorization* of size r if we can write $S = AB$ where $A \in \mathbb{R}_+^{p \times r}$ and $B \in \mathbb{R}_+^{r \times q}$ have nonnegative

entries. The smallest r for which such a factorization exists is called the *nonnegative rank* of S and denoted $\text{rank}_+(S)$.

The matrix shown in Equation (2.3), which compiles all the facets of a polytope (evaluated at the vertices) is known as the *slack matrix* of P. Formally we have:

Definition 3 (Slack matrix)**.** Let P be a polytope in $\mathbb{R}^n$. The *slack matrix* of P is a matrix where rows are indexed by vertices of P and columns are indexed by facets of P, and is defined as follows: The value at the (x, ℓ) entry of P (where x is a vertex and $\ell \leq \ell_{\max}$ is a facet inequality of P) is given by:

$$S_{x,\ell} = \ell_{\max} - \ell(x).$$

Remark 3. The slack matrix of P is not, strictly speaking, uniquely defined since it depends on the ordering of the facets/vertices and the scaling of the facet inequalities. The results stated here however will be independent of the choice of slack matrix and that is why we will often talk about "the" slack matrix of P.

It is clear that the entries of a slack matrix are nonnegative. It is not difficult to see that Theorem 1 can be equivalently written in terms of nonnegative factorizations of the slack matrix.

Theorem 2 (Yannakakis [101]; restatement of Theorem 1)**.** *Let $P = \text{conv}(X)$ be a full-dimensional polytope and let S be its slack matrix. Then P has a LP lift of size d if, and only if, S has a nonnegative factorization of size d. As a consequence, the smallest size of a LP lift of P is equal to* $\text{rank}_+(S)$.

Proof. If S has a nonnegative factorization $S = AB$ of size d then we can interpret each column of A as a nonnegative function $a_i : X \to \mathbb{R}_+$, and the condition of Theorem 1 is satisfied (the coefficients b_i's are given by the matrix B). The other direction is similar. □

2.1.4 Proof of Yannakakis' theorem

We now present a proof of Yannakakis' theorem.

Proof of Theorem 1. We first prove necessity. Assume $P = \pi(\mathbb{R}^d_+ \cap L)$ is LP lift of P of size d, where L is an affine subspace of $\mathbb{R}^d$ and π a linear map. Let $\ell_{\max} \leq \ell$ be any facet inequality of P. Since $\pi(\mathbb{R}^d_+ \cap L) \subseteq P$ it is clear that the following implication holds for any y:

$$\begin{cases} y \geq 0 \\ y \in L \end{cases} \Rightarrow \ell \circ \pi(y) \leq \ell_{\max}.$$

One can show using Farkas' lemma/strong duality for LP, that there exist coefficients $b_1, \ldots, b_d \geq 0$ and an affine form γ that vanishes on L such that

$$\ell_{\max} - \ell \circ \pi(y) = \sum_{i=1}^{d} b_i y_i + \gamma(y). \tag{2.6}$$

Note that Equation (2.6) automatically implies that $\ell_{\max} - \ell \circ \pi$ is nonnegative on $\mathbb{R}^d_+ \cap L$. To see why (2.6) is true let L_0 be the linear subspace of $\mathbb{R}^d$ parallel to L and let $y_0 \in \mathbb{R}^d$ such that $L = y_0 + L_0$. Consider the following primal/dual pair of LPs whose value is $\ell_{\max}$:

$$\begin{array}{lll} \max\limits_{y \in \mathbb{R}^d} & (\ell \circ \pi)(y) & \\ \text{s.t.} & y \in \mathbb{R}^d_+ & \\ & y - y_0 \in L_0 & \end{array} \qquad \begin{array}{lll} \min\limits_{b,h \in \mathbb{R}^d} & -\langle h, y_0 \rangle & \\ \text{s.t.} & -\ell \circ \pi = b + h & \\ & b \in \mathbb{R}^d_+, \ h \in L_0^\perp & \end{array} \tag{2.7}$$

By strong duality, there exists $b \in \mathbb{R}^d_+$ and $h \in L_0^\perp$ such that $-\ell \circ \pi = b + h$ and $-\langle h, y_0 \rangle = \ell_{\max}$. Thus this means that

$$\ell_{\max} - \ell \circ \pi = b + h - \langle h, y_0 \rangle$$

which is exactly (2.6) with $\gamma(y) = \langle h, y - y_0 \rangle$. Now for any $x \in X$ we know that there exists $A(x) = (a_1(x), \ldots, a_d(x)) \in \mathbb{R}^d_+ \cap L$ such that $\pi(A(x)) = x$. By evaluating Equation (2.6) at $y = A(x)$ we get (using the fact that $A(x) \in L$ and so $\gamma(A(x)) = 0$):

$$\ell_{\max} - \ell(x) = \sum_{i=1}^{d} b_i a_i(x) \quad \forall x \in X.$$

Thus this proves the claim.

We now show sufficiency. Assume $P = \{x \in \mathbb{R}^n : Fx \leq g\}$ is a facet description of our polytope P where $F \in \mathbb{R}^{N \times n}$ and $g \in \mathbb{R}^N$. Assume $S = AB$ is a nonnegative factorization of the slack matrix of size d, where $A \in \mathbb{R}^{|X| \times d}_+$ and $B \in \mathbb{R}^{d \times N}_+$. Then it is easy to verify that P can be written as:

$$P = \{x \in \mathbb{R}^n : \exists a \in \mathbb{R}^d \text{ s.t. } a \geq 0, \ g - Fx = B^T a\}. \tag{2.8}$$

To see why the inclusion "$\subseteq$" holds, note that if $x \in X$ is a vertex of P then by letting a be the row of A indexed by x the constraints on the right-hand side are satisfied. The inclusion "$\supseteq$" is trivial since $B^T a \geq 0$ for $a \geq 0$. The proof is almost complete since the right-hand side of (2.8) is defined using only d linear inequalities. To be sure we just need to show that (2.8) can be put in the form $P = \pi(\mathbb{R}^d_+ \cap L)$ for some linear map π and affine subspace $L \subset \mathbb{R}^d$. Since P is bounded and $\dim(P) > 0$, we know that $\text{rank}(F) = n$ and $g \notin \text{Im}(F)$. Since $\text{rank}(F) = n$ the equation $Fx = g - B^T a$ (in x) has a unique solution x_a if $g - B^T a \in \text{Im}(F)$ and no solution otherwise. It is easy to see that the map that sends a to x_a, defined on the affine subspace $L = \{a \in \mathbb{R}^d : g - B^T a \in \text{Im}(F)\}$, is affine. Since 0 does not belong to L (this is because $g \notin \text{Im}(F)$) this affine map can be extended to a linear map π on the whole space. We thus finally get that $P = \pi(\mathbb{R}^d_+ \cap L)$. This shows that P admits a LP lift of size d. $\square$

2.1.5 Pseudo-expectation point of view

In the proof of Theorem 1 we showed how to construct a lift of a polytope given certificates of nonnegativity of the facet inequalities (cf. (2.8)). In this section we give an alternative point of view of this lift. This point of view may look more abstract than (2.8) however it is more general and gives a better understanding of where the lift comes from. The same ideas will reappear when we consider positive semidefinite lifts.

The starting point of the lift is the following trivial representation of $P = \text{conv}(X)$, which follows simply from the definition of convex hull:

$$\text{conv}(X) = \left\{ \int_X x d\mu(x) : \ \mu \text{ probability measure on } X \right\}. \tag{2.9}$$

This expression says that $\text{conv}(X)$ is the set of first moments of probability measures supported on X. Consider the expectation operator E of a probability measure μ, which is given by:

$$E(f) = \int_X f(x) d\mu(x)$$

where f is any real-valued function on X. Note that Equation (2.9) can be equivalently written in terms of expectation operators as:

$$\begin{aligned} \text{conv}(X) = \big\{ (E(e_1), \ldots, E(e_n)) \ : E \text{ is the expectation operator of some} \\ \text{probability measure } \mu \text{ supported on } X \big\} \end{aligned} \tag{2.10}$$

where $e_1, \ldots, e_n$ are the coordinate functions, i.e., $e_i(x) = x_i$. The key difficulty in describing $\text{conv}(X)$ is, therefore, in understanding expectation operators of probability measures on X. It is not difficult to come up with necessary conditions for a map E to be an expectation operator. Clearly it must satisfy $E(1) = 1$ (where 1 is the constant function equal to 1), and it must also satisfy $E(a) \geq 0$ whenever a is a nonnegative function on X.

We now go back to the setting of Theorem 1 and recall that we have functions $a_1, \ldots, a_d : X \to \mathbb{R}_+$ that are nonnegative on X. Any expectation operator must thus satisfy $E(a_i) \geq 0$ for all $i = 1, \ldots, d$. Given the functions $a_1, \ldots, a_d$ we can thus construct the following *relaxation* of $\text{conv}(X)$:

$$\begin{aligned} \text{conv}(X) \subseteq \big\{ (\widetilde{E}(e_1), \ldots, \widetilde{E}(e_n)) \ : \widetilde{E} \in (\mathbb{R}^X)^* \text{ s.t. } \widetilde{E}(1) = 1, \\ \widetilde{E}(a_i) \geq 0, \ \forall i = 1, \ldots, d \big\}. \end{aligned} \tag{2.11}$$

We used the notation $\widetilde{E}$ instead of E since the maps $\widetilde{E}$ are not necessarily expectations of probability measures on X (though that is how we want to think of them). Note that $\widetilde{E}$ is an element of the dual space $(\mathbb{R}^X)^*$ since it is a linear map that takes a function on X (an element of $\mathbb{R}^X$) and outputs a real number.

Recall now that our functions a_i from Theorem 1 satisfy a very specific property. This property precisely allows us to show that (2.11) is, in fact, an equality. To

see why this is the case, let x be a point in the right-hand side of (2.11), i.e., $x = (\widetilde{E}(e_1), \ldots, \widetilde{E}(e_n))$ for some $\widetilde{E}$ that satisfies $\widetilde{E}(1) = 1$ and $\widetilde{E}(a_i) \geq 0$ for all $i = 1, \ldots, d$. We will prove that $x \in \text{conv}(X)$ by showing that $\ell(x) \leq \ell_{\max}$ for any facet inequality $\ell \leq \ell_{\max}$ of $\text{conv}(X)$. Let thus $\ell \leq \ell_{\max}$ be a facet inequality of $\text{conv}(X)$. By our assumption on the a_i's from Theorem 1, we know that there exist coefficients $b_1, \ldots, b_d \geq 0$ such that $\ell_{\max} - \ell|_X = \sum_{i=1}^d b_i a_i$. Since $x = (\widetilde{E}(e_1), \ldots, \widetilde{E}(e_n))$ and $\ell(x) = \sum_{i=1}^n \ell_i x_i$ (in the canonical basis) we have:

$$\begin{aligned}\ell_{\max} - \ell(x) = \ell_{\max} - \sum_{i=1}^n \ell_i \widetilde{E}(e_i) &\overset{(a)}{=} \widetilde{E}\left(\ell_{\max} - \sum_{i=1}^n \ell_i e_i\right) \\ &\overset{(b)}{=} \widetilde{E}\left(\sum_{i=1}^d b_i a_i\right) = \sum_{i=1}^d b_i \widetilde{E}(a_i) \geq 0\end{aligned}$$

where in (a) we used the linearity of $\widetilde{E}$ and the fact that $\widetilde{E}(1) = 1$, and in (b) we used the assumption that $\ell_{\max} - \ell|_X = \sum_{i=1}^d b_i a_i$. We have thus proved that $\ell_{\max} - \ell(x) \geq 0$. Since this is true for any facet inequality of $\text{conv}(X)$ we have thus shown that $x \in \text{conv}(X)$.

The maps $\widetilde{E}$ are known as *pseudo-expectations* (most notably in the theoretical computer science literature [2]) since they act as expectations of probability distributions even though they are not necessarily such. We will revisit pseudo-expectations in the next section when discussing SDP lifts.

2.2 SDP lifts

In this section we treat semidefinite programming (SDP) lifts of polytopes. A *spectrahedron* of size d is a convex set Q that can be described using a linear matrix inequality of size d, i.e.,

$$Q = \{y \in \mathbb{R}^N : F_0 + y_1 F_1 + \cdots + y_N F_N \in \mathbf{S}_+^d\}$$

where $F_0, \ldots, F_N$ are real symmetric matrices of size $d \times d$ and $\mathbf{S}_+^d$ denotes the cone of $d \times d$ positive semidefinite matrices. Equivalently if we call L the affine subspace of $\mathbf{S}^d$ defined as $F_0 + \text{span}(F_1, \ldots, F_N)$ we can think of a spectrahedron as the intersection of the cone $\mathbf{S}_+^d$ with this affine subspace. For convenience this is the definition we will be adopt in this thesis.

Definition 4. Let P be a polytope. We say that P has a SDP lift of size d if it can be written as $P = \pi(\mathbf{S}_+^d \cap L)$ where $\mathbf{S}_+^d$ is the cone of $d \times d$ real symmetric positive semidefinite matrices, and L is an affine subspace of $\mathbf{S}^d$. The size of the smallest SDP lift of P is called the *SDP extension complexity* of P and denoted $\text{xc}_{\text{SDP}}(P)$.

Example 2. Figure 2-4 illustrates a SDP lift of the square $[-1,1]^2$ of size 3, given by:

$$[-1,1]^2 = \left\{ (x_1,x_2) \in \mathbb{R}^2 : \exists u \in \mathbb{R} \begin{bmatrix} 1 & x_1 & x_2 \\ x_1 & 1 & u \\ x_2 & u & 1 \end{bmatrix} \succeq 0 \right\}. \tag{2.12}$$

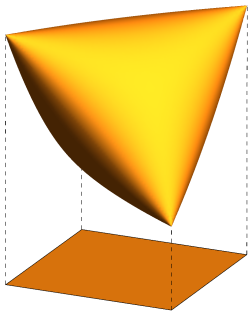

Figure 2-4: SDP lift of the square $[-1,1]^2$ of size 3 (cf. Equation (2.12)). The three-dimensional convex set shown in the figure is the set of (x_1,x_2,u) such that the 3×3 symmetric matrix on the right-hand side of (2.12) is positive semidefinite. Projecting this set onto (x_1,x_2) yields the square $[-1,1]^2$.

To see why (2.12) is true note that if $(x_1,x_2) \in \{-1,1\}^2$ then by letting $u = x_1x_2$ we have:

$$\begin{bmatrix} 1 & x_1 & x_2 \\ x_1 & 1 & u \\ x_2 & u & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ x_1 \\ x_2 \end{bmatrix} \begin{bmatrix} 1 \\ x_1 \\ x_2 \end{bmatrix}^T \succeq 0.$$

Since the right-hand side of (2.12) is convex this shows that $\text{conv}(\{-1,1\}^2) = [-1,1]^2$ is contained in it. Conversely if (x_1,x_2) belongs to the right-hand side of (2.12) then by looking at the 2×2 minors of the 3×3 positive semidefinite matrix we easily get that $1-x_1^2 \geq 0$ and $1-x_2^2 \geq 0$ i.e., $(x_1,x_2) \in [-1,1]^2$. $\lozenge$

2.2.1 Factorization theorem SDP lifts

One can prove a result similar to Theorem 1 which characterizes SDP lifts of a polytope P in terms of certificates of nonnegativity of the facet inequalities. This theorem is due to Gouveia, Parrilo, Thomas [50].

Theorem 3 (Gouveia, Parrilo, Thomas, [50])**.** *Let P be a full-dimensional polytope and let X be its set of vertices. The polytope P has a SDP lift of size d if, and only if, there exists a map $A : X \to \mathbf{S}^d_+$ such that the following holds: for any facet inequality $\ell \leq \ell_{\max}$ of P there exists $B \in \mathbf{S}^d_+$ such that*

$$\ell_{\max} - \ell(x) = \langle A(x), B \rangle \quad \forall x \in X. \tag{2.13}$$

Note that Equation (2.13) is an equality of functions on X. Just like in the LP case, it should be understood as a certificate of nonnegativity for $\ell_{\max} - \ell|_X$. Indeed the function $x \mapsto \langle A(x), B\rangle$ is "obviously" nonnegative on X since $A(x) \in \mathbf{S}^d_+$ and $B \in \mathbf{S}^d_+$.

Remark 4. If the SDP lift has the form $P = \pi(\mathbf{S}^d_+ \cap L)$ then, as we will see in the proof of the theorem, the map $A : X \to \mathbf{S}^d_+$ can be any map that satisfies $A(x) \in \mathbf{S}^d_+ \cap L$ and $\pi(A(x)) = x$ for any $x \in X$. We note this property here since it will be useful in later chapters.

Example 3. To illustrate Theorem 3, let us go back to the example of the square $[-1, 1]^2$ (Example 2) and let us exhibit the function $A : X \to \mathbf{S}^3_+$ in this case. In this example the vertex set is $X = \{-1, 1\}^2$. Consider the map A given by:

$$A(x) = \begin{bmatrix} 1 & x_1 & x_2 \\ x_1 & 1 & x_1x_2 \\ x_2 & x_1x_2 & 1 \end{bmatrix}.$$

Note that $A(x) \in \mathbf{S}^3_+$ for $x \in X$ since we have, for any $x \in X$ (using the fact that $x_1^2 = x_2^2 = 1$):

$$A(x) = \begin{bmatrix} 1 \\ x_1 \\ x_2 \end{bmatrix} \begin{bmatrix} 1 \\ x_1 \\ x_2 \end{bmatrix}^T \succeq 0.$$

To show that the condition of Theorem 3 is satisfied consider the facet inequality $1 - x_1 \geq 0$. Define

$$B = \frac{1}{2} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}^T$$

and note that $B \succeq 0$. Then we have for any $x \in \{-1, 1\}^2$

$$\langle A(x), B\rangle = \left\langle \begin{bmatrix} 1 & x_1 & x_2 \\ x_1 & 1 & x_1x_2 \\ x_2 & x_1x_2 & 1 \end{bmatrix}, \frac{1}{2} \begin{bmatrix} 1 & -1 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \right\rangle = \frac{1}{2}(1 - 2x_1 + 1) = 1 - x_1.$$

Similarly one can show that the three other facet inequalities $1 + x_1 \geq 0, 1 - x_2 \geq 0$ and $1 + x_2 \geq 0$ can be written as $\langle A(x), B\rangle$ for a suitable choice of B. This shows that the map A satisfies the condition of Theorem 3. ♢

2.2.2 Positive semidefinite factorizations

In the same way that Theorem 1 can be expressed in terms of nonnegative factorization of the slack matrix of P, Theorem 3 can similarly be formulated in terms of so-called *positive semidefinite factorizations* of a matrix.

Definition 5 (Gouveia, Parrilo, Thomas [50])**.** Let $S \in \mathbb{R}^{p \times q}_+$ be a matrix with nonnegative entries. We say that S has a *positive semidefinite factorization (psd*

factorization) of size d if there exist positive semidefinite matrices $A_1, \dots, A_p \in \mathbf{S}^d_+$ and $B_1, \dots, B_q \in \mathbf{S}^d_+$ such that $S_{ij} = \langle A_i, B_j \rangle$ for all $i = 1, \dots, p$ and $j = 1, \dots, q$. The size of the smallest psd factorization of S is called the *psd rank* of S and denoted $\operatorname{rank}_{\text{psd}}(S)$.

Theorem 3 can now be formulated in terms of positive semidefinite factorizations of the slack matrix of P (recall the definition of slack matrix, Definition 3).

Theorem 4 (Gouveia, Parrilo, Thomas [50]; restatement of Theorem 3)**.** *Let P be a full-dimensional polytope and let S be its slack matrix. Then P has a SDP lift of size d if, and only if, S has a positive semidefinite factorization of size d. As a consequence, the smallest size of a SDP lift of P is equal to* $\operatorname{rank}_{\text{psd}}(S)$.

For more information on the positive semidefinite rank, we refer the reader to the paper [32] which surveys some of its properties and applications in optimization as well as in other areas.

2.2.3 Sums of squares

In this section we show that SDP lifts can also be interpreted in terms of *sum of squares* certificates of the facet inequalities $\ell_{\max} - \ell|_X$. Such a certificate consists in expressing $\ell_{\max} - \ell|_X$ as a sum of squares of functions on X. More formally we have:

Theorem 5. *Let P be a full-dimensional polytope with vertex set X. Assume there is a subspace V of $\mathbb{R}^X$ such that the following holds:*

(*) for any facet inequality $\ell \leq \ell_{\max}$ of P there are elements $h_1, \dots, h_J \in V$ such that

$$\ell_{\max} - \ell|_X = \sum_{j=1}^{J} h_j^2. \tag{2.14}$$

Then P has a SDP lift of size $\dim V$.

Conversely if P has a SDP lift of size d, then there is a subspace V of $\mathbb{R}^X$ of dimension at most d^2 such that condition (*) holds.

Proof. We start by proving the first part. Assume we have a subspace V of dimension d such that condition (*) holds. Let $f_1, \dots, f_d$ be a basis of this subspace and define the map $A : X \to \mathbf{S}^d_+$ as follows:

$$A(x) = \begin{bmatrix} f_1(x) \\ \vdots \\ f_d(x) \end{bmatrix} \begin{bmatrix} f_1(x) \\ \vdots \\ f_d(x) \end{bmatrix}^T = [f_i(x) f_j(x)]_{1 \leq i,j \leq d}.$$

Clearly $A(x) \succeq 0$ (note also that $A(x)$ is rank-one for any $x \in X$). Now let $\ell \leq \ell_{\max}$ be a facet inequality for P. By (*) we know that there exists elements $h_1, \dots, h_J \in V$ such that (2.14) holds. We will show that there exists $B \in \mathbf{S}^d_+$ such that $\ell_{\max} - \ell|_X = \langle A(\cdot), B \rangle$. Since each h_j is in V we can write $h_j(x) = b_j^T f(x)$ where $b_j \in \mathbb{R}^d$ and

$f(x)$ is the column vector $[f_1(x); \ldots; f_d(x)]$. Now note that $(h_j(x))^2 = (b_j^T f(x))^2 = \langle f(x)f(x)^T, b_j b_j^T \rangle = \langle A(x), b_j b_j^T \rangle$. Thus if we let $B = \sum_{j=1}^{J} b_j b_j^T$ we get

$$\ell_{\max} - \ell(x) = \langle A(x), B \rangle \quad \forall x \in X$$

which is what we wanted. Thus using Theorem 3, this shows that P has an SDP lift of size d.

We now prove the other direction. Assume that P has a SDP lift of size d. According to Theorem 3 this means that there exists $A : X \to \mathbf{S}_+^d$ such that the following holds: for any facet inequality $\ell \leq \ell_{\max}$ there is $B \in \mathbf{S}_+^d$ such that $\ell_{\max} - \ell(x) = \langle A(x), B \rangle$, for all $x \in X$. Since $A(x) \in \mathbf{S}_+^d$ we can factorize it as $A(x) = R(x)R(x)^T$. Let V be the subspace of $\mathbb{R}^X$ spanned by the entries of R, i.e., $V = \operatorname{span}(x \mapsto R_{ij}(x), i, j = 1, \ldots, d) \subset \mathbb{R}^X$ and note that $\dim V \leq d^2$. Now given $\ell \leq \ell_{\max}$ a facet inequality of P we know that there exists $B \in \mathbf{S}_+^d$ such that $\ell_{\max} - \ell(x) = \langle A(x), B \rangle$ for all $x \in X$. We can write B as $B = CC^T$ to get

$$\ell_{\max} - \ell(x) = \langle R(x)R(x)^T, CC^T \rangle = \|C^T R(x)\|_F^2.$$

Since each entry of $x \mapsto C^T R(x)$ is an element of V the previous equation gives a sum-of-squares certificate of $\ell_{\max} - \ell|_X$ using functions from V. This completes the proof. □

Example 4. The lift of the square $[-1, 1]^2$ discussed earlier can be explained in terms of sum-of-squares certificates. Note that the facet inequality $1 - x_1 \geq 0$ can be certified using sum-of-squares as follows:

$$1 - x_1 = \frac{1}{2}(1 - x_1)^2 \quad \forall x \in \{-1, 1\}^2.$$

It is crucial to note that the equality above is understood on $\{-1, 1\}^2$ (the equality is of course not true globally because the left hand side is a polynomial of degree 1 whereas the right-hand side is a polynomial of degree 2). Since we are working on $\{-1, 1\}^2$ the right-hand side expands to $\frac{1}{2}(1 - 2x_1 + x_1^2) = \frac{1}{2}(1 - 2x_1 + 1) = 1 - x_1$ where we used the fact that $x_1^2 = 1$. Similarly one can show that the other facet inequalities have the following sum-of-squares certificates:

$$\begin{aligned}
1 - x_2 &= \frac{1}{2}(1 - x_2)^2 \quad \forall x \in \{-1, 1\}^2, \\
1 + x_1 &= \frac{1}{2}(1 + x_1)^2 \quad \forall x \in \{-1, 1\}^2, \\
1 + x_2 &= \frac{1}{2}(1 + x_2)^2 \quad \forall x \in \{-1, 1\}^2.
\end{aligned}$$

Thus if we let V be the space of polynomials of degree at most 1 on $\{-1, 1\}^2$ (i.e., by an abuse of notation $V = \operatorname{span}(1, x_1, x_2)$) the condition (*) of Theorem 5 holds true. Thus this shows that $[-1, 1]^2$ has a SDP lift of size $\dim V = 3$. ◊

2.2.4 Proof of the factorization theorem for SDP lifts

We now give a proof of Theorem 3.

Proof of Theorem 3. The proof follows the same steps as the proof of Theorem 1 for LP lifts. The only difference is that we use the generalized Farkas' lemma/strong duality for SDPs.

Assume that $P = \pi(\mathbf{S}^d_+ \cap L)$ is an SDP lift of P of size d where L is an affine subspace of $\mathbf{S}^d$ and π a linear map. Let $\ell \leq \ell_{\max}$ be any facet inequality of P. Since $\pi(\mathbf{S}^d_+ \cap L) \subseteq P$ the following implication holds for all $Y \in \mathbf{S}^d$:

$$\begin{cases} Y \succeq 0 \\ Y \in L \end{cases} \Rightarrow \ell_{\max} - \ell \circ \pi(Y) \geq 0.$$

We will now show using Farkas' lemma/strong duality for SDP that there is a positive semidefinite matrix $B \in \mathbf{S}^d_+$, an affine form γ that vanishes on L such that

$$\ell_{\max} - \ell \circ \pi(Y) = \langle B, Y \rangle + \gamma(Y) \quad \forall Y \in \mathbf{S}^d. \tag{2.15}$$

To see why (2.15) is true let L_0 be the linear space in $\mathbf{S}^d$ parallel to L and let $Y_0 \in \mathbf{S}^d$ such that $L = Y_0 + L_0$. The following problems are dual to each other and the value of the primal (maximization) problem is equal to $\ell_{\max}$:

$$\begin{array}{lll} \max\limits_{Y \in \mathbf{S}^d} & (\ell \circ \pi)(Y) & \\ \text{s.t.} & Y \in \mathbf{S}^d_+ & \\ & Y - Y_0 \in L_0 & \end{array} \qquad \begin{array}{lll} \min\limits_{B,H \in \mathbf{S}^d} & -\langle H, Y_0 \rangle & \\ \text{s.t.} & -\ell \circ \pi = B + H & \\ & B \in \mathbf{S}^d_+, \ H \in L_0^\perp & \end{array} \tag{2.16}$$

We can assume that the intersection of L with the interior of $\mathbf{S}^d_+$ is nonempty (otherwise the intersection lies on a strict face of $\mathbf{S}^d_+$ which means that one can reduce the size of the SDP lift). In this case strong duality holds, the optimal values of the two SDPs (2.16) are equal to $\ell_{\max}$ and the dual (minimization) problem is attained. Let B, H be the optimal points of the dual problem in (2.16). From dual feasibility we have $-\ell \circ \pi = B + H$ and so since $\ell_{\max} = -\langle H, Y_0 \rangle$ we get that:

$$\ell_{\max} - \ell \circ \pi = B + H - \langle H, Y_0 \rangle.$$

Note that this shows (2.15) where the affine map γ is $\gamma(Y) = \langle H, Y - Y_0 \rangle$. For $x \in X$ let $A(x)$ be any element in $\mathbf{S}^d_+ \cap L$ such that $\pi(A(x)) = x$. Evaluating (2.15) at $A(x)$, for any $x \in X$ we get:

$$\ell_{\max} - \ell(x) = \langle B, A(x) \rangle$$

where we used the fact that $\pi(A(x)) = x$ and that $\gamma(A(x)) = 0$ since $A(x) \in L$. This proves our claim.

We now prove the converse. We show how to construct an SDP lift of P from a psd factorization of its slack matrix. Assume $P = \{x \in \mathbb{R}^n : Fx \leq g\}$ is a facet description of P where $F \in \mathbb{R}^{N \times n}$, $g \in \mathbb{R}^N$. Let S be the slack matrix of P and let

$S_{x,\ell} = \langle A(x), B(\ell)\rangle$ be a psd factorization of S of size d, where $A(x), B(\ell) \in \mathbf{S}^d_+$ (here $x \in X$ is a vertex of P and $\ell \leq \ell_{\max}$ is a facet-defining inequality of P). It is easy to verify that we have the following description of P:

$$P = \left\{x \in \mathbb{R}^n : \exists A \in \mathbf{S}^d, A \succeq 0 \text{ and } g_j - f_j^T x = \langle A, B_j\rangle \ \forall j = 1, \ldots, N\right\} \tag{2.17}$$

where f_j^T is the j'th row of F. To see why the inclusion "$\subseteq$" holds let $x \in X$ be a vertex of P and take $A = A(x)$ in the right-hand side. The reverse inclusion "$\supseteq$" follows immediately by observing that $\langle A, B_j\rangle \geq 0$ since A and B_j are positive semidefinite.

The remaining part of the proof is to show that (2.17) is indeed a positive semidefinite lift of P, i.e., that it can be put in the form $P = \pi(\mathbf{S}^d_+ \cap L)$ for some linear map π and affine subspace L. Let $T : \mathbf{S}^d \to \mathbb{R}^N$ be the linear map defined by $T(A) = (\langle A, B_1\rangle, \ldots, \langle A, B_N\rangle)$. Then we can rewrite (2.17) as:

$$P = \left\{x \in \mathbb{R}^n : \exists A \in \mathbf{S}^d_+ \text{ s.t. } g - Fx = T(A)\right\}.$$

Since P is bounded and $\dim(P) > 0$, we know that $\operatorname{rank}(F) = n$ and $g \notin \operatorname{Im}(F)$. Since $\operatorname{rank}(F) = n$ the equation $Fx = g - T(A)$ (in x) has a unique solution x_A if $g - T(A) \in \operatorname{Im}(F)$ and no solution otherwise. It is easy to see that the map that sends A to x_A, defined on the affine subspace $L = \{A : g - T(A) \in \operatorname{Im}(F)\}$, is affine. Since 0 does not belong to L (this is because $g \notin \operatorname{Im}(F)$) this affine map can be extended to a linear map π on the whole space. We thus finally get that $P = \pi(\mathbf{S}^d_+ \cap L)$. This shows that P is the projection of a spectrahedron of size d. □

2.2.5 Pseudo-expectation point of view

We now give the pseudo-expectation point of view of the lift (2.17) as we did in the case of LP lifts (cf. Section 2.1.5). Recall that $\operatorname{conv}(X)$ has the following trivial representation which simply comes from the definition of the convex hull:

$$\operatorname{conv}(X) = \left\{(E(e_1), \ldots, E(e_n)) \ : \begin{array}{r} E \text{ is the expectation operator of some} \\ \text{probability measure } \mu \text{ supported on } X\end{array}\right\}. \tag{2.18}$$

Assume now that we have a map $A : X \to \mathbf{S}^d_+$ as in Theorem 3. Note that if E is any valid expectation operator on X then it has to satisfy $E_x(A(x)) \in \mathbf{S}^d_+$ (the subscript x is just to indicate that we are taking the expectation with respect to x). By simply imposing this condition and the normalization constraint $E(1) = 1$ we thus get the following relaxation of $\operatorname{conv}(X)$:

$$\operatorname{conv}(X) \subseteq \left\{(\widetilde{E}(e_1), \ldots, \widetilde{E}(e_n)) \ : \ \widetilde{E} \in (\mathbb{R}^X)^*, \widetilde{E}(1) = 1, \widetilde{E}_x(A(x)) \in \mathbf{S}^d_+\right\}. \tag{2.19}$$

In the same way as in the LP case, one can show that if the condition of Theorem 3 is satisfied then we have equality in (2.19). We omit the proof here since it is very similar to the LP case explained in Section 2.1.5.

The case of sum-of-squares lifts (see Section 2.2.3) also has a simple interpretation in this setting. It simply consists in enforcing the constraint that $E(f^2) \geq 0$ for all $f \in V$, where V is the subspace of functions in Theorem 5. Under the conditions of Theorem 5 we can show:

$$\operatorname{conv}(X) = \left\{ (\widetilde{E}(e_1), \ldots, \widetilde{E}(e_n)) : \widetilde{E} \in (\mathbb{R}^X)^*, \widetilde{E}(1) = 1, \widetilde{E}(f^2) \geq 0 \; \forall f \in V \right\}. \quad (2.20)$$

Note that the constraint $\widetilde{E}(f^2) \geq 0$ for all $f \in V$ can be written as a positive semidefinite constraint of size $\dim V$ since it expresses the fact that the quadratic form $f \in V \mapsto \widetilde{E}(f^2)$ is positive semidefinite. The symmetric matrix associated to this quadratic form is often called a *moment matrix*.

As we mentioned earlier when discussing the LP case the map $\widetilde{E}$ is often called a *pseudo-expectation*. In fact this terminology is most often used in the case of the sum-of-squares relaxations, see e.g., [2].

2.3 Hierarchies

The theorems presented in the previous sections show that constructing a lift of a polytope P is equivalent to finding certificates of nonnegativity for the facet inequalities. The question is: how do we find such certificates? how can we find, in a systematic way, the functions $a_1, \ldots, a_d$ of Theorem 1 (LP lifts), or the subspace V of Theorem 5 (sum-of-squares lifts)? There is of course no magical way of producing such functions in general but some of the existing *hierarchies* can be shown to correspond to specific choices. In this section we briefly outline these choices.

2.3.1 Krivine/Handelman/Sherali-Adams hierarchy

Assume that our polytope $P = \operatorname{conv}(X)$ is 0-1, i.e., that the vertex set X is a subset of $\{0,1\}^n$, and more precisely that it can be written as

$$X = \{x \in \{0,1\}^n : g_1(x) \geq 0, \ldots, g_m(x) \geq 0\}$$

where $g_1, \ldots, g_m$ are some polynomials. In this case it is not difficult to come up with functions that are nonnegative on X. In fact for any choice of subsets $I \subseteq T \subseteq [n]$ and integers $\gamma \in \mathbb{N}^m$ the following function is nonnegative on X:

$$a_{T,I,\gamma}(x) = \prod_{i \in I} x_i \prod_{i \in T \setminus I} (1 - x_i) \prod_{j=1}^{m} g_j(x)^{\gamma_j}.$$

Note that $a_{T,I,\gamma}$ is a polynomial of degree at most $|T| + \sum_{j=1}^m \gamma_j \deg(g_j)$. For a fixed integer k we can consider the outer-relaxation (2.11) of P obtained by considering only the functions $a_{T,I,\gamma}$ of degree at most k. By increasing k we get a hierarchy of increasingly tighter relaxations to our polytope $P = \operatorname{conv}(X)$. Certificates of this form are the basis of the well-known Sherali-Adams hierarchy [96], and also of the

Handelman hierarchy [58, 75]. In fact such certificates have been investigated as early as 1964 by Krivine in [66]. We refer the reader to [71, 75] for more details on the specifics of each hierarchy (which functions to include at the level k of the hierarchy) and for questions related to the convergence of the hierarchy.

2.3.2 Lasserre/theta-body hierarchy

One of the most studied methods to produce SDP lifts of polytopes is the so-called Lasserre/theta-body hierarchy [69, 49]. This method can be explained very simply in terms of the terminology set up in Section 2.2.5. The relaxation at level k is exactly given by (2.20) where the subspace V consists of the space of polynomials of degree at most k. More explicitly, the k'th level of the Lasserre/theta-body hierarchy for $\text{conv}(X)$ can be expressed as:

$$\begin{aligned}\text{TH}_k(X) := \big\{(\widetilde{E}(e_1), \ldots, \widetilde{E}(e_n)) \ : \ & \widetilde{E} \in (\mathbb{R}^X)^*, \widetilde{E}(1) = 1 \\ & \widetilde{E}(f^2) \geq 0 \ \forall f \in \text{Pol}_{\leq k}(X)\big\}\end{aligned} \tag{2.21}$$

where $\text{Pol}_{\leq k}(X)$ is the space of polynomials of degree at most k on $X \subset \mathbb{R}^n$, i.e., it is the restriction to X of polynomials in $\mathbb{R}[x_1, \ldots, x_n]$ of degree at most k:

$$\text{Pol}_{\leq k}(X) := \{f \in \mathbb{R}^X : \exists p \in \mathbb{R}[x_1, \ldots, x_n]_{\leq k} \text{ s.t. } p(x) = f(x) \ \forall x \in X\}.$$

The notation TH in (2.21) is for "theta-body", see [49]. The smallest k such that $\text{TH}_k(X) = \text{conv}(X)$ is known as the *theta-rank* of X. The Lasserre/theta-body relaxations have been extensively studied in combinatorial optimization and theoretical computer science, as well as in the more recent field of convex algebraic geometry; we refer the reader to [10, 74, 2] for more details.

2.4 Complexity-theoretic considerations

Before concluding this chapter we discuss in this section some complexity-theoretic implications related to the existence/inexistence of polynomial-size LP/SDP lifts for polytopes arising from combinatorial optimization problems.

Many combinatorial optimization problems can be formulated using linear programming over a "naturally"-defined polytope. Consider for example the *traveling salesman problem* which asks to find the minimum weight Hamiltonian cycle on a given weighted graph G. We can assume for simplicity that the graph G is the complete graph on n nodes (this does not affect the computational complexity of the problem). To model this problem using linear programming, define $\chi_S \in \mathbb{R}^E$ to be the characteristic vector of a subset $S \subseteq E$ of the edges of the complete graph:

$$\chi_S(e) = \begin{cases} 1 & \text{if } e \in S \\ 0 & \text{else.} \end{cases}$$

The TSP polytope is defined as the convex hull of all characteristic vectors of Hamiltonian cycles in the complete graph K_n:

$$\mathrm{TSP}(n) := \mathrm{conv}\left\{\chi_S : S \text{ Hamiltonian cycle in } K_n\right\} \subset \mathbb{R}^{\binom{n}{2}}.$$

Given a weight function $w : E \to \mathbb{R}_+$ the minimum weight Hamiltonian cycle can be obtained by solving the following linear program:

$$\text{minimize} \quad \sum_{e \in E} w(e)x(e) \quad \text{subject to} \quad x \in \mathrm{TSP}(n).$$

One can similarly define the *matching polytope*, the *stable set* polytope, the *cut* polytope which are associated to their respective combinatorial optimization problems. For example the (perfect) matching polytope on a complete n-node graph is defined as:

$$\mathrm{MATCH}(n) := \mathrm{conv}\left\{\chi_S : S \text{ perfect matching in } K_n\right\} \subset \mathbb{R}^E.$$

It is clear that if one can find a LP lift for the TSP polytope of polynomial-size then one could solve the traveling salesman problem in polynomial-time using e.g., the ellipsoid method or path-following methods. With some additional technical details the same would also be true if we have a polynomial-size SDP lift [2].

It was shown by Fiorini et al. [41] that the TSP polytope does not, in fact, have a polynomial-size LP lift, and more recently Lee et al. [76] generalized their result to show that it does not admit a polynomial-size SDP lift. It is important to note however that these results do not imply in any way that TSP is not in the complexity class **P**.[3] The fact that $\mathrm{TSP}(n)$ does not admit a polynomial-size LP lift does not rule out for example that $\mathrm{TSP}(n)$ could have a polynomial-time separation oracle: in fact it was shown recently by Rothvoß [89] that the matching polytope does not admit a polynomial-size LP lift despite having a well-known efficient separation oracle [57]. Also note that when expressing the TSP problem using linear programming we used a certain encoding of the problem in terms of the characteristic vectors, that is admittedly natural, but nevertheless not the only one. A different encoding of the problem could yield a different polytope with different extension complexities.

[2] More precisely, in addition to having $\mathrm{TSP}(n) = \pi_n(\mathbf{S}_+^{d(n)} \cap L_n)$ where $d(n)$ is polynomial in n and a description of π_n and L_n that can be generated in time polynomial n, we also need to have $X_n \in \mathbf{S}^{d(n)}$, $r_n, R_n > 0$ such that $B(X_n, r_n) \subseteq \mathbf{S}_+^{d(n)} \cap L_n \subseteq B(X_n, R_n)$ where $\log(R_n/r_n) = \mathrm{poly}(n)$ and where $B(x, r)$ is the Euclidean ball centered at x with radius r. Under these assumptions one can get, using the ellipsoid method or interior-point method, an ϵ-approximation of the TSP problem in time polynomial in the problem size and $\log(1/\epsilon)$, for any $\epsilon > 0$. We refer the reader to [57] and [28] for more details on the complexity results concerning SDP using ellipsoid method and interior-point methods.

[3] There was some confusion about this in some online discussions, see e.g., the comments section of http://mat.tepper.cmu.edu/blog/?p=1587, https://spokutta.wordpress.com/2012/01/05/1311/, http://blog.computationalcomplexity.org/2014/04/favorite-theorems-extended-formulations.html.

2.5 Summary of chapter

We summarize briefly the main results in this section.

- Let $P = \mathrm{conv}(X)$ be a polytope. Any facet inequality $\ell(x) \leq \ell_{\max}$ where $\ell_{\max} := \max_{x \in X} \ell(x)$ can be seen as a nonnegative function $\ell_{\max} - \ell$ on X. Constructing a small lift of P is equivalent to finding "small" certificates of nonnegativity of all the facet inequalities $\ell_{\max} - \ell|_X$.

- For LP lifts, the certificates of nonnegativity that we want have the form

$$\ell_{\max} - \ell(x) = \sum_{i=1}^{d} b_i a_i(x) \quad \forall x \in X$$

  where $a_1, \ldots, a_d$ are fixed nonnegative functions on X (i.e., independent of the facet $\ell \leq \ell_{\max}$) and $b_1, \ldots, b_d \geq 0$ depend on ℓ. The size of the lift in this case is d.

- For SDP lifts, the certificates of nonnegativity take the form

$$\ell_{\max} - \ell(x) = \langle A(x), B \rangle \quad \forall x \in X$$

  where $A : X \to \mathbf{S}^d_+$ is fixed and $B \in \mathbf{S}^d_+$ depends on ℓ.

- Sum-of-squares lifts (which are a special case of SDP lifts) consist in finding a subspace V of $\mathbb{R}^X$ such that any facet inequality has a certificate:

$$\ell_{\max} - \ell|_X = \sum_{j=1}^{J} h_j^2$$

  where $h_1, \ldots, h_J \in V$. Such a subspace yields an SDP lift of P of size $\dim V$.

Chapter 3

Nonnegative rank

In this chapter we study the nonnegative rank which was introduced in Chapter 2 (see Definition 2). After briefly reviewing existing techniques to obtain lower bounds on the nonnegative rank we propose a new method to obtain such bounds which unifies some of the existing techniques and which inherits many of the structural properties of the nonnegative rank (invariance under scaling, subadditivity, monotonicity, etc.). Our technique also applies to other special notions of rank and we outline such extensions briefly at the end of the chapter. The content of this chapter is based on the paper [34].

3.1 Preliminaries

Recall from Chapter 2, Definition 2 that the nonnegative rank of an entrywise nonnegative matrix $A \in \mathbb{R}_+^{m \times n}$ is the smallest r such that we can find $U \in \mathbb{R}_+^{m \times r}$ and $V \in \mathbb{R}_+^{r \times n}$ such that $A = UV$. Equivalently, $\text{rank}_+(A)$ can also be defined as the smallest r such that there exist nonnegative rank-one matrices $R_1, \ldots, R_r$ such that

$$A = \sum_{i=1}^{r} R_i. \tag{3.1}$$

The correspondence between a decomposition (3.1) and a factorization $A = UV$ is $R_i = u_i v_i^T$ where u_i is the i'th column of U and v_i^T the i'th row of V. Note that $\text{rank}_+(A)$ always satisfies:

$$\text{rank}(A) \leq \text{rank}_+(A) \leq \min(m, n).$$

Applications of nonnegative rank We saw in Chapter 2, Theorem 2 that the nonnegative rank of the slack matrix of a polytope P is equal to the size of the smallest LP lift of P. The nonnegative rank also plays an important role in other areas such as statistical modeling [30, 67] and communication complexity [78, 81]. In statistical modeling the matrix A is interpreted as the joint probability distribution of a pair of random variables (X, Y):

$$A(x, y) = \mathbb{P}[X = x, Y = y].$$

In this context a nonnegative factorization of A of size r consists in expressing (X, Y) as a mixture of r pairs of random variables (X_i, Y_i) where X_i and Y_i are independent. Indeed such a decomposition takes the form:

$$\mathbb{P}[X = x, Y = y] = \sum_{i=1}^{r} \mathbb{P}[W = i] \cdot \mathbb{P}[X = x | W = i] \cdot \mathbb{P}[Y = y | W = i],$$

where W is the mixing distribution, taking values in $\{1, \ldots, r\}$ and X and Y are conditionally independent given W (the pair (X_i, Y_i) is given by the conditional distribution $(X, Y)|W = i$). This is exactly a nonnegative factorization of A of the form (3.1) where the rank-one matrix R_i is given by $R_i(x, y) = \mathbb{P}[W = i]\mathbb{P}[X = x|W = i]\mathbb{P}[Y = y|W = i]$.

Another application of the nonnegative rank is in communication complexity where one is interested in the minimum number of bits that need to be exchanged between two parties in order to compute a binary function $f : \mathcal{X} \times \mathcal{Y} \to \{0, 1\}$, assuming that initially each party holds only one of the two arguments of the function. This quantity is known as the communication complexity of f and is tightly related to the nonnegative rank of the $|\mathcal{X}| \times |\mathcal{Y}|$ matrix M_f associated to f defined by $M_f(x, y) = f(x, y)$ [78, 81].

Notations Throughout this chapter a matrix M is called nonnegative if it is entry-wise nonnegative. For convenience we define the following partial order on the indices of a matrix $A \in \mathbb{R}^{m \times n}$ that will be used later:

$$(i,j) \leq (k,l) \;\Leftrightarrow\; i \leq k \text{ and } j \leq l, \tag{3.2}$$

and we write $(i,j) < (k,l)$ if $i < k$ and $j < l$.

3.2 Existing methods to lower bound the nonnegative rank

In this section we briefly review some of the existing techniques to bound the nonnegative rank. We start by reviewing combinatorial lower bounds on the nonnegative rank which only use the support pattern of the matrix A. We then describe so-called hyperplane separation bounds, and finally we discuss bounds based on information theoretic quantities.

3.2.1 Combinatorial bounds

Let $A \in \mathbb{R}_+^{m \times n}$ be a nonnegative matrix. A *monochromatic rectangle* for A is a rectangle $R = I \times J \subseteq [m] \times [n]$ such that $A_{i,j} > 0$ for any $(i,j) \in R$, i.e., the rectangle does not touch any zero entry of A.

Definition 6. The *boolean rank* of A (also called the *rectangle covering number*), denoted $\text{rank}_B(A)$ is the minimum number of monochromatic rectangles needed to cover all the nonzero entries of A.

Note that in any nonnegative factorization $A = \sum_{i=1}^{r} u_i v_i^T$, the rectangles $R_i = \text{supp}(u_i) \times \text{supp}(v_i)$ are necessarily monochromatic for A. From this observation we get that:

$$\text{rank}_B(A) \leq \text{rank}_+(A).$$

Example 5. Let $A = I_n$ be the identity matrix of size n. It is easy to see that any monochromatic rectangle for I_n is a singleton and has the form $R = \{i\} \times \{i\}$ where $i = 1, \ldots, n$. Thus $\text{rank}_B(I_n) = n$. ◊

Rectangle graph As noted in [40] the boolean rank of A can be expressed as the chromatic number of a certain graph constructed from A. Define the *rectangle graph* of A, denoted $\text{RG}(A)$ as follows: the vertex set of $\text{RG}(A)$ is the set of indices (i,j) such that $A_{i,j} > 0$; furthermore there is an edge (undirected) between vertices (i,j) and (k,l) if, and only if, $A_{i,l}A_{k,j} = 0$. Note that if two entries (i,j) and (k,l) of A are connected by an edge in $\text{RG}(A)$, then the two entries cannot be covered by the same monochromatic rectangle (see Figure 3-1 for an illustration).

$$\begin{bmatrix} \boldsymbol{A_{i,j}} & A_{i,l} \\ A_{k,j} & \boldsymbol{A_{k,l}} \end{bmatrix}$$

Figure 3-1: If $A_{i,j} > 0$ and $A_{k,l} > 0$ and one of $A_{i,l}$ or $A_{k,j}$ is zero, then it is not possible to cover $A_{i,j}$ and $A_{k,l}$ with the same monochromatic rectangle. In this case we put an edge between vertices (i, j) and (k, l) in the graph $\mathrm{RG}(A)$.

Using this observation, it is not hard to show that the minimum number of monochromatic rectangles needed to cover the nonzero entries A is precisely the chromatic number of $\mathrm{RG}(A)$ [40, Lemma 5.3]:

$$\mathrm{rank}_B(A) = \chi(\mathrm{RG}(A)).$$

An obvious lower bound on the chromatic number of $\mathrm{RG}(A)$ is the *clique number* of $\mathrm{RG}(A)$, i.e., the size of the largest clique, which is denoted by $\omega(\mathrm{RG}(A))$. The clique number $\omega(\mathrm{RG}(A))$ is also sometimes known as the *fooling set number* of A. Other well-known lower bounds on $\chi(\mathrm{RG}(A))$ are the *fractional chromatic number* $\chi_{\mathrm{frac}}(\mathrm{RG}(A))$ and the *(complement) Lovász theta number* $\overline{\vartheta}(\mathrm{RG}(A))$. These quantities satisfy the following inequalities:

$$\boxed{\mathrm{fool}(A) = \omega(\mathrm{RG}(A)) \leq \overline{\vartheta}(\mathrm{RG}(A)) \leq \chi_{\mathrm{frac}}(\mathrm{RG}(A)) \leq \chi(\mathrm{RG}(A)) = \mathrm{rank}_B(A)}.$$

Application Fiorini et al. showed in [41] that the cut polytope (among others) does not admit polynomial-sized extended formulation. In order to do so, they showed that the boolean rank of (a submatrix of) the slack matrix of the cut polytope is superpolynomial.

3.2.2 Hyperplane separation bounds

Another bounding technique that proved powerful for the nonnegative rank is the so-called hyperplane separation bound. This technique was used by Rothvoß in his major result [89] where he obtained an exponential lower bound on the LP extension complexity of the matching polytope. In the next proposition we denote by $\|A\|_\infty = \max_{ij} |A_{ij}|$ the entrywise infinity norm of a matrix A.

Proposition 1 (Hyperplane separation bound for the $\|\cdot\|_\infty$ norm)**.** *Let $A \in \mathbb{R}_+^{m \times n}$ be a nonnegative matrix. Assume that $L : \mathbb{R}^{m \times n} \to \mathbb{R}$ is a linear map that satisfies the following assumption:*

$$L(R) \leq 1 \text{ for any nonnegative rank-one matrix } R \in \mathbb{R}_+^{m \times n} \text{ satisfying } \|R\|_\infty = 1. \tag{3.3}$$

Then we have

$$\mathrm{rank}_+(A) \geq \frac{L(A)}{\|A\|_\infty}. \tag{3.4}$$

Proof. Let $A = \sum_{i=1}^{r} R_i$ be a nonnegative decomposition of A with $r = \operatorname{rank}_+(A)$ terms. Since all the R_is are elementwise nonnegative we have, for each $i = 1, \ldots, r$, the elementwise inequalities $0 \leq R_i \leq A$. Using this observation, the bound (3.4) then follows easily from the following sequence of inequalities:

$$L(A) = \sum_{i=1}^{r} L(R_i) \overset{(a)}{\leq} \sum_{i=1}^{r} \|R_i\|_\infty \overset{(b)}{\leq} \sum_{i=1}^{r} \|A\|_\infty = r\|A\|_\infty$$

where in (a) we used the fact that $L(R_i) = \|R_i\|_\infty L(\frac{1}{\|R_i\|_\infty} R_i) \leq \|R_i\|_\infty$ which follows from the hypothesis (3.3) and the fact that the entrywise maximum of $R_i/\|R_i\|_\infty$ is 1, and in (b) we used the fact that $0 \leq R_i \leq A$ which implies that $\|R_i\|_\infty \leq \|A\|_\infty$. □

As it is clear from the proof of Proposition 1, there is nothing specific about the infinity norm, except monotonicity, which makes the bounding technique possible. We discuss generalizations of Proposition 1 with other norms in more detail in Section 3.3.6.

3.2.3 Information-theoretic bounds

Information theoretic quantities can also be used to get lower bounds on the nonnegative rank; in fact such bounds were used in [17, 18, 16] in the context of extended formulations of polytopes. To see how these lower bounds work, recall from Section 3.1 that if A is a nonnegative matrix representing the joint distribution of a pair of random variables (X, Y), then a nonnegative factorization of A of size r expresses the fact that (X, Y) is a mixture of r independent random variables. Using this interpretation, the nonnegative rank of A can thus be formulated as:

$$\operatorname{rank}_+(A) = \min_{\substack{X-W-Y \\ (X,Y)\sim A}} |\operatorname{supp}(W)|, \tag{3.5}$$

where $|\operatorname{supp}(W)|$ is the number of values that W takes, and where the Markov chain constraint $X-W-Y$ means that X and Y are conditionally independent given W (the random variable W plays the role of the mixing distribution). The formulation (3.5) allows us to draw connections between the nonnegative rank and certain information-theoretic quantities. Consider for example the Wyner common information $C(X;Y)$ defined in [100] as:

$$C(X;Y) = \min_{X-W-Y} I(XY;W).$$

Wyner showed in [100] that $C(X;Y)$ quantifies the minimum number of bits that two parties need to share in order to generate samples from the joint distribution (X, Y). Since $I(XY;W) \leq \log|\operatorname{supp}(W)|$ it follows that $C(X;Y)$ is always a lower bound for $\log \operatorname{rank}_+(A)$:

$$C(X;Y) \leq \log \operatorname{rank}_+(A). \tag{3.6}$$

The lower bound (3.6) has been used in the context of extended formulations of polytopes and we refer the reader to [17, 16] for more details.

3.3 Self-scaled bounds for nonnegative rank

In this section we introduce our new bound on the nonnegative rank and outline its connection to combinatorial as well as hyperplane separation bounds. The method we propose can be applied to a general class of ranks but we will mainly focus here on the nonnegative rank. We mention at the end (Section 3.3.7) the possible extensions to other notions of ranks.

3.3.1 Definition

We start by explaining the main idea of the lower bound. Let $A \in \mathbb{R}_+^{m\times n}$ and consider a decomposition of A of the form:

$$A = \sum_{i=1}^{r} R_i \tag{3.7}$$

where R_i, for $i = 1, \dots, r$ are rank-one and nonnegative. An important observation is that each term R_i in the decomposition above necessarily satisfies

$$0 \leq R_i \leq A$$

where $\leq$ denotes entrywise inequality of matrices. In other words, if we define the set:

$$\mathcal{A}_+(A) := \Big\{ R \in \mathbb{R}^{m\times n} \ : \ \operatorname{rank} R \leq 1 \text{ and } 0 \leq R \leq A \Big\}, \tag{3.8}$$

then in any decomposition of A of the form (3.7), all the terms R_i must necessarily belong to $\mathcal{A}_+(A)$. As a consequence, if we can produce a linear functional L such that $L(R) \leq 1$ for all $R \in \mathcal{A}_+(A)$, then clearly $L(A)$ is a lower bound on the minimal number of terms in any decomposition of A of the form (3.7). Indeed this is because we have:

$$L(A) = \sum_{i=1}^{r} L(R_i) \leq \sum_{i=1}^{r} 1 = r.$$

Thus for such an L we have $L(A) \leq \operatorname{rank}_+(A)$. Now to obtain the best lower bound, one can look for the linear functional L which maximizes the value of $L(A)$ while satisfying $L \leq 1$ on $\mathcal{A}_+(A)$. We call this quantity $\tau_+(A)$ and this is the main object we study in this section:

$$\tau_+(A) := \max_{L \text{ linear}} \ L(A) \quad \text{subject to} \quad L(R) \leq 1 \ \ \forall R \in \mathcal{A}_+(A). \tag{3.9}$$

3.3.2 Duality and self-scaled property

Minimization formulation of τ_+ Using convex duality, one can obtain a dual formulation of $\tau_+(A)$ as the solution of a certain minimization problem. In fact the next lemma shows that $\tau_+(A)$ is nothing but the *atomic norm* of A [22] associated to the set of atoms $\mathcal{A}_+(A)$. This interpretation of $\tau_+(A)$ will be very useful later when

studying its properties.

Lemma 1. *If $A \in \mathbb{R}_+^{m\times n}$ then we have:*

$$\tau_+(A) = \min\{t > 0 \ : \ A \in t\,\mathrm{conv}(\mathcal{A}_+(A))\}. \tag{3.10}$$

In other words, $\tau_+(A)$ is the Minkowski gauge function of $\mathrm{conv}(\mathcal{A}_+(A))$*, evaluated at A.*

Proof. Observe that Equation (3.9) expresses the fact that $\tau_+(A)$ is the support function of $\mathrm{conv}(\mathcal{A}_+(A))^\circ$, evaluated at A. Theorem 14.5 in [88] shows that the support function of the polar C° of a closed convex set C is equal to the Minkowski gauge function of C. Thus it follows that $\tau_+(A)$ is equal to the Minkowski gauge function of $\mathrm{conv}(\mathcal{A}_+(A))$, evaluated at A, which is precisely Equation (3.10). □

Illustration The next example illustrates the geometric picture underlying the atomic norm formulation of $\tau_+(A)$.

Example 6. Assume A is a 2×2 diagonal matrix $A = \mathrm{diag}(a_1, a_2)$ where $a_i \geq 0$. In this case one can easily verify that $\mathcal{A}_+(A)$ is given by:

$$\begin{aligned}\mathcal{A}_+(A) &= \left\{ R \in \mathbb{R}^{2\times 2} \ : \ \mathrm{rank}\, R \leq 1 \text{ and } 0 \leq R \leq \begin{bmatrix} a_1 & 0 \\ 0 & a_2 \end{bmatrix} \right\} \\ &= \left\{ \begin{bmatrix} x & 0 \\ 0 & 0 \end{bmatrix} \text{ with } 0 \leq x \leq a_1 \right\} \cup \left\{ \begin{bmatrix} 0 & 0 \\ 0 & y \end{bmatrix} \text{ with } 0 \leq y \leq a_2 \right\}.\end{aligned} \tag{3.11}$$

The convex hull of $\mathcal{A}_+(A)$ (projected onto the diagonal elements) is depicted in Figure 3-2. Observe that, when $a_1, a_2 > 0$, the smallest t such that $A \in t\,\mathrm{conv}(\mathcal{A}_+(A))$ is $t = 2$ and thus $\tau_+(A) = 2 = \mathrm{rank}_+(A)$. ◊

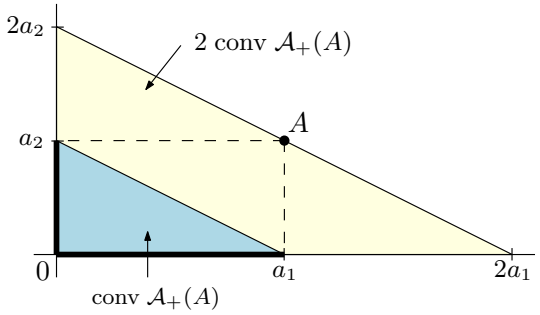

Figure 3-2: Depiction of the set of atoms $\mathcal{A}_+(A)$ and its convex hull for a 2×2 diagonal matrix A (cf. Equation (3.11)). The set $\mathcal{A}_+(A)$ consists of the two dark heavy lines joining the origin to a_1 and a_2. The convex hull of $\mathcal{A}_+(A)$ is formed by the triangle $0, a_1, a_2$.

Self-scaled property We can see in Example 6 the *self-scaled*[1] feature of the bound $\tau_+(A)$. This is in contrast with the existing hyperplane separation methods to lower bound $\text{rank}_+(A)$ where the scaling of the atoms is *independent* of A: for example in Proposition 1 the scaling is done using the entrywise infinity norm, and in [33] the scaling is with respect to the Frobenius norm. This feature is explained in more detail in Section 3.3.6 where we show that τ_+ always yields better bounds than any such hyperplane separation method.

3.3.3 Semidefinite relaxation

The quantity $\tau_+(A)$ defined in (3.9) cannot be efficiently computed in general, since we do not have an efficient description of the feasible set $\{L \text{ linear } : \ L(R) \leq 1 \ \forall R \in \mathcal{A}_+(A)\}$ (note however that (3.9) is a convex optimization problem). In this section we introduce a semidefinite programming relaxation of $\tau_+(A)$. To do so, we construct an over-relaxation of the set $\text{conv}(\mathcal{A}_+(A))$ which can be represented using linear matrix inequalities. Recall that $\mathcal{A}_+(A)$ is the intersection of the variety of rank-one matrices with the set $\{R \in \mathbb{R}^{m\times n} \ : \ 0 \leq R \leq A\}$. The variety of rank-one matrices is described by the vanishing of 2×2 minors, i.e.,

$$R_{i,j}R_{k,l} - R_{i,l}R_{k,j} = 0 \tag{3.12}$$

for all $(1,1) \leq (i,j) < (k,l) \leq (m,n)$ (recall the partial order $(i,j) < (k,l) \Leftrightarrow i < k$ and $j < l$, see Equation (3.2)). Let $r = \text{vec}(R)$ be the vector obtained by stacking all the columns of R and consider the following positive-semidefinite matrix:

$$\begin{bmatrix} 1 \\ r \end{bmatrix} \begin{bmatrix} 1 \\ r \end{bmatrix}^T = \begin{bmatrix} 1 & r^T \\ r & rr^T \end{bmatrix}. \tag{3.13}$$

Note that rr^T is a symmetric $mn \times mn$ matrix whose rows and columns are indexed by entries of R. The quadratic equations (3.12) corresponding to the vanishing of 2×2 minors of R can be written as linear equations in the entries of rr^T, namely:

$$(rr^T)_{ij,kl} - (rr^T)_{il,kj} = 0$$

for $(1,1) \leq (i,j) < (k,l) \leq (m,n)$ (in the equation above, the subscripts "ij" and "kl" in $(rr^T)_{ij,kl}$ are the indices in $\{1,\ldots,mn\}$ for the entries (i,j) and (k,l) respectively—we will use this slight abuse of notation to avoid having heavy notations).
Also note that the inequality $R \leq A$ implies that:

$$(rr^T)_{ij,ij} \leq r_{ij}A_{ij} \tag{3.14}$$

[1] We use the word *self-scaled* as a descriptive term to convey the main idea of the lower bound presented here. It is not related to the term as used in the context of interior-point methods (e.g., "self-scaled barrier" [84]).

which is a linear inequality in the entries of the matrix (3.13). Using these two observations we have the following over-relaxation of $\mathrm{conv}(\mathcal{A}_+(A))$:

$$\mathrm{conv}(\mathcal{A}_+(A)) \subseteq \mathcal{A}_+^{\mathrm{sos}}(A) \tag{3.15}$$

where

$$\begin{aligned}\mathcal{A}_+^{\mathrm{sos}}(A) = \Bigg\{ R \in \mathbb{R}^{m\times n} \ : \exists X \in \mathbf{S}^{mn} \ \text{ such that } \ & \begin{bmatrix} 1 & \mathrm{vec}(R)^T \\ \mathrm{vec}(R) & X \end{bmatrix} \succeq 0 \\ & \text{and } X_{ij,ij} \leq R_{ij}A_{ij} \quad \forall i \in [m], j \in [n] \\ & \text{and } X_{ij,kl} - X_{il,kj} = 0 \quad \forall (1,1) \leq (i,j) < (k,l) \leq (m,n) \Bigg\}.\end{aligned} \tag{3.16}$$

If we define $\tau_+^{\mathrm{sos}}(A)$ as:

$$\tau_+^{\mathrm{sos}}(A) = \min\{t > 0 \ : \ A \in t\mathcal{A}_+^{\mathrm{sos}}(A)\}$$

then we clearly have (by the inclusion (3.15)):

$$\tau_+^{\mathrm{sos}}(A) \leq \tau_+(A) \leq \mathrm{rank}_+(A).$$

Furthermore, the quantity $\tau_+^{\mathrm{sos}}(A)$ can be computed using semidefinite programming. Indeed, it is not difficult to show using the description (3.16) of $\mathcal{A}_+^{\mathrm{sos}}(A)$ that we have:

$$\begin{aligned}\tau_+^{\mathrm{sos}}(A) \quad = \quad \min_{t,X} \quad & t \\ \text{s.t.} \quad & \begin{bmatrix} t & \mathrm{vec}(A)^T \\ \mathrm{vec}(A) & X \end{bmatrix} \succeq 0 \\ & X_{ij,ij} \leq A_{ij}^2 \quad \forall i \in [m], j \in [n] \\ & X_{ij,kl} = X_{il,kj} \quad \forall (1,1) \leq (i,j) < (k,l) \leq (m,n)\end{aligned} \tag{3.17}$$

Duality and sum-of-squares interpretation The dual of the semidefinite program (3.17) takes the form of a *sum-of-squares* program, namely we have[2]:

$$\begin{aligned}\tau_+^{\mathrm{sos}}(A) = \quad \max \quad & L(A) \\ \text{s.t.} \quad & L \text{ is a linear form} \\ & 1 - L(X) = SOS(X) + \textstyle\sum_{ij} D_{ij}X_{ij}(A_{ij} - X_{ij}) \mod I \\ & D_{ij} \geq 0 \\ & SOS(X) \text{ is a sum-of-squares polynomial}\end{aligned} \tag{3.18}$$

[2] To obtain (3.18), we write the Lagrangian dual of (3.17) and then do the change of variables $L_{ij} := -2\Lambda_{ij} - D_{ij}A_{ij}$, where D_{ij} is the dual variable for the constraint $X_{ij,ij} \leq A_{ij}^2$ and Λ is the top-right $1 \times mn$ block of the dual variable for the constraint $\left[\begin{smallmatrix} t & \mathrm{vec}(A)^T \\ \mathrm{vec}(A) & X \end{smallmatrix}\right] \succeq 0$. The L_{ij} are the coordinates of the linear form L, i.e., $L(X) = \sum_{ij} L_{ij}X_{ij}$.

Here I is the ideal in $\mathbb{R}[X_{11}, \ldots, X_{mn}]$ corresponding to the variety of $m \times n$ rank-one matrices, i.e., it is ideal generated by the 2×2 minors $X_{ij}X_{kl} - X_{il}X_{kj}$. The sum-of-squares constraint in (3.18) means that the polynomials on each side of the equality are equal when X is rank-one. Note that this sum-of-squares constraint can be rewritten more explicitly as requiring that:

$$1 - L(X) - \sum_{ij} D_{ij}X_{ij}(A_{ij} - X_{ij}) - \sum_{(i,j)<(k,l)} \nu_{ijkl}(X_{ij}X_{kl} - X_{il}X_{kj}) \text{ is a sum-of-squares}$$

where the parameters ν_{ijkl} are real numbers[3]. It is clear that any such L satisfies $L(X) \leq 1$ for all $X \in \mathcal{A}_+(A)$. As such, (3.18) is a natural sum-of-squares relaxation of (3.9).

Zero entries in A When the matrix A has some entries equal to 0, the semidefinite program (3.17) that defines $\tau_+^{\text{sos}}(A)$ can be reduced, since in this case the feasible set is contained in a low-dimensional face of the positive semidefinite cone (such a reduction is called *facial reduction*, see e.g., [13] and [87] for more information and applications of facial reduction). Let $S = \text{supp}(A) = \{(i,j) \ : \ A_{i,j} > 0\}$ be the set of nonzero entries of A, and define

$$\pi : \mathbb{R}^{m \times n} \to \mathbb{R}^S, \quad \pi(A) = (A_{i,j})_{ij \in S}$$

to be the linear map that projects onto the entries in S. Observe that, in the SDP (3.17), if $A_{i,j} = 0$ for some (i,j) then necessarily $X_{ij,ij} = 0$. Thus by the positive semidefiniteness constraint this implies that the ij'th row and ij'th column of X are identically zero, and one can thus eliminate this row and column from the program. Using this fact, one can show that $\tau_+^{\text{sos}}(A)$ can be computed using the following reduced semidefinite program where the size of the matrix X is now $|\text{supp}(A)| \times |\text{supp}(A)|$, instead of $mn \times mn$ (recall that $\pi(A)$ is the vectorization of A where we only keep the nonzero entries of A):

$$\begin{aligned}
\tau_+^{\text{sos}}(A) \ = \ \min_{t,X} \ & t \\
\text{s.t.} \ & \begin{bmatrix} t & \pi(A)^T \\ \pi(A) & X \end{bmatrix} \succeq 0 \\
& \forall (i,j) \text{ s.t. } A_{i,j} > 0 : \ X_{ij,ij} \leq A_{ij}^2 \\
& \forall (1,1) \leq (i,j) < (k,l) \leq (m,n) \text{ s.t. } A_{i,j}A_{k,l} > 0 \text{ or } A_{i,l}A_{k,j} > 0 : \\
& \qquad \begin{cases} \text{if } A_{i,l}A_{k,j} = 0 \ : \ X_{ij,kl} = 0 \\ \text{if } A_{i,j}A_{k,l} = 0 \ : \ X_{il,kj} = 0 \\ \text{else} \quad X_{ij,kl} - X_{il,kj} = 0 \end{cases}
\end{aligned} \tag{3.19}$$

[3]One can show that the sum-of-squares polynomial cannot have degree more than 2 and the multipliers ν_{ijkl} are necessarily real numbers.

3.3.4 Structural properties

In this section we explore some of the properties of $\tau_+(A)$ and $\tau_+^{\text{sos}}(A)$. We show that $\tau_+(A)$ and $\tau_+^{\text{sos}}(A)$ have many appealing properties (invariance under diagonal scaling, subadditivity, monotonicity, etc.) which are not present in most of currently existing bounds on the nonnegative rank.

Theorem 6. *Let $A \in \mathbb{R}_+^{m\times n}$ be a nonnegative matrix.*

1. *Invariance under diagonal scaling: If D_1 and D_2 are diagonal matrices with strictly positive entries on the diagonal, then $\tau_+(D_1AD_2) = \tau_+(A)$ and $\tau_+^{\text{sos}}(D_1AD_2) = \tau_+^{\text{sos}}(A)$.*

2. *Invariance under permutation of rows or columns: If P_1 and P_2 are permutation matrices of size $m \times m$ and $n \times n$ respectively, then $\tau_+(P_1AP_2) = \tau_+(A)$ and $\tau_+^{\text{sos}}(P_1AP_2) = \tau_+^{\text{sos}}(A)$.*

3. *Subadditivity: If $B \in \mathbb{R}_+^{m\times n}$ is a nonnegative matrix then:*

$$\tau_+(A+B) \leq \tau_+(A) + \tau_+(B) \quad \text{and} \quad \tau_+^{\text{sos}}(A+B) \leq \tau_+^{\text{sos}}(A) + \tau_+^{\text{sos}}(B).$$

4. *Product: If $B \in \mathbb{R}_+^{n\times p}$, then*

$$\tau_+(AB) \leq \min(\tau_+(A), \tau_+(B)) \quad \text{and} \quad \tau_+^{\text{sos}}(AB) \leq \min(\tau_+^{\text{sos}}(A), \tau_+^{\text{sos}}(B)).$$

5. *Monotonicity: If B is a submatrix of A (i.e., $B = A[I, J]$ for some $I \subseteq [m]$ and $J \subseteq [n]$), then $\tau_+(B) \leq \tau_+(A)$ and $\tau_+^{\text{sos}}(B) \leq \tau_+^{\text{sos}}(A)$.*

6. *Block-diagonal composition: Let $B \in \mathbb{R}_+^{m'\times n'}$ be another nonnegative matrix and define*

$$A \oplus B = \begin{bmatrix} A & 0 \\ 0 & B \end{bmatrix}.$$

Then

$$\tau_+(A \oplus B) = \tau_+(A) + \tau_+(B) \quad \text{and} \quad \tau_+^{\text{sos}}(A \oplus B) = \tau_+^{\text{sos}}(A) + \tau_+^{\text{sos}}(B)$$

Proof. See Section 3.5.1. □

3.3.5 Comparison with combinatorial bounds

We will now see that the quantities $\tau_+(A)$ and $\tau_+^{\text{sos}}(A)$ can be interpreted as non-combinatorial equivalents of $\chi_{\text{frac}}(\text{RG}(A))$ and $\overline{\vartheta}(\text{RG}(A))$ respectively, where $\chi_{\text{frac}}(\text{RG}(A))$ is the fractional rectangle covering number of A and $\overline{\vartheta}(\text{RG}(A))$ is the complement of the Lovász ϑ number for the rectangle graph of A (these quantities were defined in Section 3.2.1). In fact one can prove the following theorem:

Theorem 7. *If $A \in \mathbb{R}_+^{m\times n}$ is a nonnegative matrix, then*

$$\tau_+(A) \geq \chi_{\text{frac}}(\text{RG}(A)) \quad \text{and} \quad \tau_+^{\text{sos}}(A) \geq \overline{\vartheta}(\text{RG}(A)).$$

Proof. See Section 3.5.2. □

To give insights into the relation between $\tau_+(A)$ and $\tau_+^{\text{sos}}(A)$ with $\chi_{\text{frac}}(\text{RG}(A))$ and $\overline{\vartheta}(\text{RG}(A))$ we recall below the definitions of the fractional chromatic number and the Lovász theta number and we give their expression for the rectangle graph $\text{RG}(A)$.

- The fractional chromatic number of a graph G is a linear programming relaxation of the chromatic number (note however that the size of this LP relaxation may have exponential size and the fractional chromatic number is actually NP-hard [82]). When applied to the rectangle graph of A, the quantity is called the *fractional rectangle cover* of A (see e.g., [65]). Let $\mathcal{A}_B(A)$ be the set of monochromatic rectangles valid for A (the subscript "B" here stands for "Boolean"):

  $$\mathcal{A}_B(A) = \Big\{ R \in \{0,1\}^{m\times n} \ : \ R \text{ is a monochromatic rectangle for } A \Big\}.$$

  Using this notation, the fractional rectangle cover number of A is the solution of the following linear program:

  $$\begin{aligned} \chi_{\text{frac}}(\text{RG}(A)) = \ \min \ & \sum_{R\in\mathcal{A}_B(A)} x_R \\ \text{s.t.} \ & \forall R \in \mathcal{A}_B(A) \ : \ x_R \geq 0 \\ & \forall (i,j), \ \ A_{i,j} > 0 \ \Rightarrow \sum_{R\in\mathcal{A}_B(A)} x_R R_{i,j} \geq 1. \end{aligned} \tag{3.20}$$

  Note that if we replace the constraint $x_R \geq 0$ with the binary constraint $x_R \in \{0,1\}$, we get the exact rectangle cover number of A. We can rewrite the linear program above in the following form, which emphasizes the connection with the quantity $\tau_+(A)$ (cf. Equation (3.10)):

  $$\begin{aligned} \chi_{\text{frac}}(\text{RG}(A)) = \ \min \ & t \\ \text{s.t.} \ & \exists Y \in t\,\text{conv}(\mathcal{A}_B(A)) \ \text{ s.t. } \ \forall (i,j), \ \ A_{i,j} > 0 \ \Rightarrow \ Y_{i,j} \geq 1. \end{aligned} \tag{3.21}$$

  The variable Y above plays the role of $\sum_{R\in\mathcal{A}_B(A)} x_R R$ in (3.20).

  Note that a result of Lovász [80] shows that for any graph $G = (V,E)$ the fractional chromatic number of G is always within a $\ln|V|$ factor from $\chi(G)$, namely:

  $$\frac{1}{1+\ln|V|}\chi(G) \leq \chi_{\text{frac}}(G) \leq \chi(G).$$

- Given a graph $G = (V,E)$, the complement Lovász theta number $\overline{\vartheta}(G) \stackrel{\text{def}}{=} \vartheta(\overline{G})$

is defined by the following semidefinite program:

$$\begin{aligned} \overline{\vartheta}(G) = \quad & \min && t \\ & \text{subject to} && \begin{bmatrix} t & 1^T \\ 1 & X \end{bmatrix} \succeq 0 \\ & && X_{u,u} = 1 \quad \forall u \in V \\ & && X_{u,v} = 0 \quad \forall \{u,v\} \in E \end{aligned}$$

When applied to the rectangle graph $\mathrm{RG}(A)$ of a nonnegative matrix A, we get:

$$\begin{aligned} \overline{\vartheta}(\mathrm{RG}(A)) = \quad & \min && t \\ & \text{subject to} && \begin{bmatrix} t & 1^T \\ 1 & X \end{bmatrix} \succeq 0 \\ & && \forall (i,j) \text{ s.t. } A_{i,j} > 0 : \; X_{ij,ij} = 1 \\ & && \forall (1,1) \leq (i,j) < (k,l) \leq (m,n) : \\ & && \quad \begin{cases} \text{if } A_{i,l}A_{k,j} = 0 \; : \; X_{ij,kl} = 0 & (3.22\text{a}) \\ \text{if } A_{i,j}A_{k,l} = 0 \; : \; X_{il,kj} = 0 & (3.22\text{b}) \end{cases} \end{aligned} \tag{3.22}$$

Note how the semidefinite program above resembles the semidefinite program (3.19) which defines $\tau_+^{\mathrm{sos}}(A)$.

Figure 3-3 summarizes the different quantities discussed in this section and how they relate to the quantities $\tau_+(A)$ and $\tau_+^{\mathrm{sos}}(A)$:

$$\begin{array}{ccccccc} & & \tau_+^{\mathrm{sos}}(A) & \leq & \tau_+(A) & \leq & \mathrm{rank}_+(A) \\ & & | \vee & & | \vee & & | \vee \\ \mathrm{fool}(A) = \omega(\mathrm{RG}(A)) & \leq & \overline{\vartheta}(\mathrm{RG}(A)) & \leq & \chi_{\mathrm{frac}}(\mathrm{RG}(A)) & \leq & \chi(\mathrm{RG}(A)) = \mathrm{rank}_B(A) \end{array}$$

Figure 3-3: Summary of the relations between $\tau_+(A), \tau_+^{\mathrm{sos}}(A)$ and some combinatorial lower bounds on $\mathrm{rank}_+(A)$.

3.3.6 Comparison with hyperplane separation bounds

In this section we compare our bound to hyperplane separation bounds which we saw in Section 3.2.2. We first introduce a generalization of Proposition 1 which holds for a larger class of norms (not just the entrywise infinity norm).

Definition 7. A function $\mathcal{N} : \mathbb{R}_+^{m \times n} \to \mathbb{R}_+$ is called *positively homogeneous* if it satisfies $\mathcal{N}(\lambda A) = \lambda \mathcal{N}(A)$ for any $A \in \mathbb{R}_+^{m \times n}$ and $\lambda \geq 0$. Furthermore, it is called *monotone* if for any $A, B \in \mathbb{R}_+^{m \times n}$ such that $A \leq B$ (componentwise inequality) we have $\mathcal{N}(A) \leq \mathcal{N}(B)$.

Norms on $\mathbb{R}^{m\times n}$ form a natural class of positively homogeneous functions. Many norms also satisfy the monotonicity property, like for example, the Frobenius norm (i.e., the ℓ_2 entrywise norm):

$$\|A\|_F = \sqrt{\sum_{i=1}^{m}\sum_{j=1}^{n} A_{i,j}^2},$$

or the ℓ_∞ entrywise norm:

$$\|A\|_\infty = \max_{\substack{1\le i\le m\\ 1\le j\le n}} |A_{i,j}|.$$

Define $\mathcal{A}_\mathcal{N}$ to be the set of rank-one matrices in the "unit ball" of $\mathcal{N}$, i.e.,

$$\mathcal{A}_\mathcal{N} := \{X \in \mathbb{R}_+^{m\times n} \ : \ \operatorname{rank} X \le 1 \text{ and } \mathcal{N}(X) \le 1\}. \tag{3.23}$$

We can also define:

$$\begin{aligned}\mathcal{N}^*(A) &= \min\{t > 0 \ : \ A \in t\operatorname{conv}(\mathcal{A}_\mathcal{N})\}\\ &= \max\{L(A) \ : \ L \text{ linear and } L(X) \le 1 \ \forall X \in \mathcal{A}_\mathcal{N}\}.\end{aligned} \tag{3.24}$$

The fact that the two formulations of $\mathcal{N}^*(A)$ in Equation (3.24) are equal follows from convex duality and the same arguments used in Lemma 1. The following proposition, which generalizes Proposition 1, shows that one can obtain a lower bound on $\operatorname{rank}_+(A)$ using $\mathcal{N}^*(A)$ and $\mathcal{N}(A)$:

Proposition 2. *Let $\mathcal{N} : \mathbb{R}_+^{m\times n} \to \mathbb{R}_+$ be a monotone positively homogeneous function, and let $\mathcal{N}^*$ be defined as in Equation (3.24). Then for any $A \in \mathbb{R}_+^{m\times n}$, we have:*

$$\operatorname{rank}_+(A) \ge \frac{\mathcal{N}^*(A)}{\mathcal{N}(A)}.$$

Proof. Let $A = \sum_{i=1}^r A_i$ be a decomposition of A with $r = \operatorname{rank}_+(A)$ terms and where each A_i is rank-one and nonnegative. Let L be the optimal solution in the maximization problem of Equation (3.24). Then we have:

$$\begin{aligned}\mathcal{N}^*(A) = L(A) = \sum_{i=1}^r L(A_i) &= \sum_{i=1}^r \mathcal{N}(A_i) L\left(\frac{1}{\mathcal{N}(A_i)} A_i\right)\\ &\overset{(a)}{\le} \sum_{i=1}^r \mathcal{N}(A_i) \overset{(b)}{\le} \sum_{i=1}^r \mathcal{N}(A) = r\mathcal{N}(A)\end{aligned}$$

where in (a) we used the homogeneity of $\mathcal{N}$ and the fact that $L(X) \le 1$ when $\mathcal{N}(X) \le 1$, and in (b) we used the fact that for each i we have $A_i \le A$, and thus by monotonicity of $\mathcal{N}$ we have $\mathcal{N}(A_i) \le \mathcal{N}(A)$. Thus we finally get that

$$r \ge \frac{\mathcal{N}^*(A)}{\mathcal{N}(A)}$$

which is what we wanted. □

In [33] we studied the case where $\mathcal{N}$ is the Frobenius norm, and where the associated quantity $\mathcal{N}^*$ was called the *nonnegative nuclear norm* and was denoted by ν_+. For this particular choice of $\mathcal{N}$ the following stronger lower bound was shown to hold:

$$\operatorname{rank}_+(A) \geq \left(\frac{\nu_+(A)}{\|A\|_F}\right)^2.$$

In the next theorem we show that any lower bound on rank_+ obtained from monotone positively homogeneous functions like in Proposition 2 is always dominated by $\tau_+(A)$.

Theorem 8. *Let $\mathcal{N} : \mathbb{R}_+^{m\times n} \to \mathbb{R}_+$ be a monotone positively homogeneous function, and let $\mathcal{N}^*$ be as defined in Equation (3.24). Then for any $A \in \mathbb{R}_+^{m\times n}$ we have:*

$$\operatorname{rank}_+(A) \geq \tau_+(A) \geq \frac{\mathcal{N}^*(A)}{\mathcal{N}(A)}.$$

Proof. First note that we have the inclusion

$$\frac{1}{\mathcal{N}(A)}\mathcal{A}_+(A) \subseteq \mathcal{A}_{\mathcal{N}}. \tag{3.25}$$

Indeed if R is rank-one and satisfies $0 \leq R \leq A$ then we have, by homogeneity and monotonicity of $\mathcal{N}$,

$$\mathcal{N}\left(\frac{1}{\mathcal{N}(A)}R\right) = \frac{1}{\mathcal{N}(A)}\mathcal{N}(R) \leq \frac{1}{\mathcal{N}(A)}\mathcal{N}(A) \leq 1.$$

Let L be the optimal linear form in the definition of $\mathcal{N}^*(A)$ in (3.24). Since $L \leq 1$ on $\mathcal{A}_{\mathcal{N}}$, by the inclusion (3.25) we have that $L \leq 1$ on $\frac{1}{\mathcal{N}(A)}\mathcal{A}_+(A)$ or equivalently that $\frac{1}{\mathcal{N}(A)}L \leq 1$ on $\mathcal{A}_+(A)$. Thus by definition of $\tau_+(A)$ we have

$$\tau_+(A) \geq \frac{1}{\mathcal{N}(A)}L(A) = \frac{\mathcal{N}^*(A)}{\mathcal{N}(A)}.$$

□

We now show that the quantity $\tau_+(A)$ actually fits in the class of lower bounds of Proposition 2, where the homogeneous function $\mathcal{N}$ *depends on* A. Specifically if A is a nonnegative matrix, we can define $\mathcal{N}_A$ as follows:

$$\mathcal{N}_A(X) = \min\{t > 0 \ : \ X \leq tA\}.$$

Clearly $\mathcal{N}_A$ is a monotone positively homogeneous function, and it satisfies $\mathcal{N}_A(A) = 1$. Note that the set of atoms $\mathcal{A}_{\mathcal{N}_A}$ associated to $\mathcal{N}_A$ (cf. Equation (3.23)) is nothing but

$\mathcal{A}_+(A)$. Thus it follows directly from the definition (3.24) of $\mathcal{N}^*(A)$ that $\mathcal{N}^*_A(A) = \tau_+(A)$. To summarize we can write that:

$$\tau_+(A) = \sup_{\substack{\mathcal{N} \text{ monotone and} \\ \text{positively homogeneous}}} \frac{\mathcal{N}^*(A)}{\mathcal{N}(A)}.$$

3.3.7 Extension to other notions of rank

The technique presented in this chapter can be applied not only to the nonnegative rank, but more generally to a general class of rank which we call *atomic cone ranks*.

Definition 8. Let K be a convex pointed[4] cone and V be a given algebraic variety in some Euclidean space. Given $A \in K$ we define $\text{rank}_{K,V}(A)$ to be the smallest integer r for which we can write

$$A = \sum_{i=1}^{r} R_i$$

where each $R_i \in K \cap V$. The function $\text{rank}_{K,V}$ is called the *atomic rank function* associated to K and V.

The nonnegative rank corresponds to the special case where K is the cone of nonnegative matrices in $\mathbb{R}^{m\times n}$, and V is the variety of rank-one matrices. Another example of atomic cone rank is the so-called *completely positive rank*: A symmetric matrix $A \in \mathbf{S}^n$ is called *completely-positive* [7] if it admits a decomposition of the form:

$$A = \sum_{i=1}^{r} u_i u_i^T,$$

where the vectors u_i are nonnegative. The *cp-rank* of A is defined as the smallest r for which such a decomposition of A exists. It corresponds to the atomic rank where K is the cone of completely positive matrices, and V is the variety of rank-one matrices.

To generalize the technique described in this chapter to atomic cone ranks, observe that if $A \in K$ admits a decomposition of the form:

$$A = \sum_{i=1}^{r} R_i \quad \text{where} \quad R_i \in V \cap K \ \ \forall i = 1, \ldots, r \tag{3.26}$$

then necessarily each term R_i satisfies

$$0 \preceq_K R_i \preceq_K A$$

where $\preceq_K$ denotes the inequality induced by the cone K (recall that $x \preceq_K y \Leftrightarrow y - x \in K$). We can now define the set

$$\mathcal{A}_{K,V}(A) := \Big\{ R \in V \ \text{ such that } \ 0 \preceq_K R \preceq_K A \Big\}, \tag{3.27}$$

[4] A cone K is called *pointed* if $K \cap (-K) = \{0\}$. We require that the cone K is pointed so that the order $\preceq_K$ associated to K is a valid partial order (in particular, that it is antisymmetric).

which plays the same role as $\mathcal{A}_+(A)$ defined in (3.8) in the case of nonnegative rank. Now if we can produce a linear functional L such that $L(R) \leq 1$ for all $R \in \mathcal{A}_{K,V}(A)$, then clearly $L(A)$ is a lower bound on the minimal number of terms in any decomposition of A of the form (3.26), i.e., on $\text{rank}_{K,V}(A)$. Again this is because

$$L(A) = \sum_{i=1}^{r} L(R_i) \leq \sum_{i=1}^{r} 1 = r.$$

The quantity $\tau_{K,V}(A)$ is the best lower bound on $\text{rank}_{K,V}(A)$ one can obtain this way:

$$\tau_{K,V}(A) := \max_{L \text{ linear}} \quad L(A) \quad \text{subject to} \quad L(R) \leq 1 \quad \forall R \in \mathcal{A}_{K,V}(A). \tag{3.28}$$

We refer the reader to the paper [34] for more examples where one can use this bounding technique, and for the specific cases of the nonnegative tensor rank and the cp-rank.

3.4 Summary of chapter

- Let $A \in \mathbb{R}^{m \times n}$ be a nonnegative matrix. Define $\mathcal{A}_+(A)$ to be the set of nonnegative rank-one matrices that are entrywise smaller than A:

  $$\mathcal{A}_+(A) := \Big\{ R \in \mathbb{R}^{m \times n} \ : \ \text{rank}\, R \leq 1 \text{ and } 0 \leq R \leq A \Big\}.$$

  Then the quantity

  $$\tau_+(A) = \min\{t > 0 \ : \ A \in t\, \text{conv}(\mathcal{A}_+(A))\}.$$

  satisfies $\tau_+(A) \leq \text{rank}_+(A)$.

- $\tau_+(A)$ can be interpreted as a non-combinatorial version of the fractional rectangle cover number of A, see Figure 3-3. The sum-of-squares relaxation $\tau_+^{\text{sos}}(A)$ of $\tau_+(A)$ is also related to the Lovász theta number of the rectangle graph of A.

- $\tau_+(A)$ is the best bound one can get using the *hyperplane separation* technique (where "best" is in the sense of the best norm used), see Theorem 8.

- $\tau_+(A)$ and $\tau_+^{\text{sos}}(A)$ inherit many of the structural properties of the nonnegative rank (invariance under scaling, subadditivity, monotonicity, ...), see Theorem 6.

- A similar approach can be used to obtain bounds on other notions of atomic cone ranks, such as the cp-rank [34].

3.5 Proofs

3.5.1 Proof of Theorem 6 on the structural properties of τ_+ and $\tau_+^{\mathbf{sos}}$

Invariance under diagonal scaling

1. We first prove invariance under diagonal scaling for τ_+, then we consider τ_+^{sos}. Let $A' = D_1 A D_2$ where D_1 and D_2 are two diagonal matrices with strictly positive entries on the diagonal. Observe that the set of atoms $\mathcal{A}_+(A')$ of A' can be obtained from the atoms $\mathcal{A}_+(A)$ of A as follows:

$$\mathcal{A}_+(A') = \{D_1 R D_2 \ : \ R \in \mathcal{A}_+(A)\} =: D_1 \mathcal{A}_+(A) D_2. \tag{3.29}$$

Indeed, if R is rank-one and $0 \leq R \leq A$ then clearly $D_1 R D_2$ is rank-one and satisfies $0 \leq D_1 R D_2 \leq D_1 A D_2 = A'$ thus $D_1 R D_2 \in \mathcal{A}_+(A')$. Conversely if $R' \in \mathcal{A}_+(A')$, then by letting $R = D_1^{-1} R D_2^{-1}$ we see that $R' = D_1 R D_2$ with R rank-one and $0 \leq R \leq A$. Thus this shows equality (3.29). Thus we have:

$$\begin{aligned}
\tau_+(A') &= \min\{t \ : \ A' \in t \operatorname{conv}(\mathcal{A}_+(A'))\} \\
&= \min\{t \ : \ D_1 A D_2 \in t \operatorname{conv}(D_1 \mathcal{A}_+(A) D_2)\} \\
&= \min\{t \ : \ D_1 A D_2 \in t D_1 \operatorname{conv}(\mathcal{A}_+(A)) D_2\} \\
&= \min\{t \ : \ A \in t \operatorname{conv}(\mathcal{A}_+(A))\} \\
&= \tau_+(A).
\end{aligned} \tag{3.30}$$

2. We now prove invariance under diagonal scaling for the SDP relaxation τ_+^{sos}. For this we use the maximization formulation of τ_+^{sos} given in Equation (3.18). Let L be the optimal linear form in (3.18) for the matrix A, i.e., $L(A) = \tau_+^{\text{sos}}(A)$ and L satisfies:

$$1 - L(X) = SOS(X) + \sum_{ij} D_{ij} X_{ij} (A_{ij} - X_{ij}) \bmod I. \tag{3.31}$$

Define the linear polynomial $L'(X) = L(D_1^{-1} X D_2^{-1})$. We can verify that:

$$\begin{aligned}
1 - L'(X) &= 1 - L(D_1^{-1} X D_2^{-1}) \\
&= SOS(D_1^{-1} X D_2^{-1}) + \sum_{ij} D_{ij} \frac{X_{ij}}{(D_1)_{ii}(D_2)_{jj}} \left(A_{ij} - \frac{X_{ij}}{(D_1)_{ii}(D_2)_{jj}} \right) \bmod I \\
&= SOS(D_1^{-1} X D_2^{-1}) + \sum_{ij} D_{ij} \frac{X_{ij}}{(D_1)_{ii}^2 (D_2)_{jj}^2} (A'_{ij} - X_{ij}) \bmod I
\end{aligned}$$

where in the last equality we used the fact that $A_{ij} = \frac{A'_{ij}}{(D_1)_{ii}(D_2)_{jj}}$. Thus this shows that L' is feasible for the sum-of-squares program (3.18) for the matrix A'. Thus since $L'(A') = L(A) = \tau_+^{\text{sos}}(A)$, we get that $\tau_+^{\text{sos}}(A') \geq \tau_+^{\text{sos}}(A)$. With

the same reasoning we can show that:

$$\tau_+^{\text{sos}}(A) = \tau_+^{\text{sos}}(D_1^{-1}(D_1AD_2)D_2^{-1}) \geq \tau_+^{\text{sos}}(D_1AD_2) = \tau_+^{\text{sos}}(A').$$

Thus we have $\tau_+^{\text{sos}}(A') = \tau_+^{\text{sos}}(A)$.

Invariance under permutation

The proof of invariance under permutation is very similar to the one for invariance under diagonal scaling. To prove the claim for τ_+ one proceeds by showing that the set of atoms $\mathcal{A}_+(A')$ of $A' = P_1AP_2$ can be obtained from the atoms of A by applying the permutations P_1 and P_2, namely:

$$\mathcal{A}_+(A') = \{P_1RP_2 \ : \ R \in \mathcal{A}_+(A)\} =: P_1\mathcal{A}_+(A)P_2.$$

For the SDP relaxation we also use the same idea as the previous proof by constructing a certificate L' for A' using the certificate L for A. We omit the details here since they are very similar to the previous proof.

Subadditivity

1. We first prove the subadditivity property for τ_+, i.e., $\tau_+(A+B) \leq \tau_+(A)+\tau_+(B)$. Observe that we have

$$\mathcal{A}_+(A) \cup \mathcal{A}_+(B) \subseteq \mathcal{A}_+(A+B). \tag{3.32}$$

   Indeed if $R \in \mathcal{A}_+(A)$, i.e., R is rank-one and $0 \leq R \leq A$, then we also have $0 \leq R \leq A+B$ (since B is nonnegative) and thus $R \in \mathcal{A}_+(A+B)$. Thus this shows $\mathcal{A}_+(A) \subseteq \mathcal{A}_+(A+B)$, and the same reason gives $\mathcal{A}_+(B) \subseteq \mathcal{A}_+(A+B)$, and thus we get (3.32). By definition of $\tau_+(A)$ and $\tau_+(B)$, we know there exist decompositions of A and B of the form $A = \sum_i \alpha_i R_i$ and $B = \sum_j \beta_j S_j$ where $\sum_i \alpha_i = \tau_+(A)$ and $\sum_j \beta_j = \tau_+(B)$ and where $R_i \in \mathcal{A}_+(A)$ for all i and $S_j \in \mathcal{A}_+(B)$ for all j. Adding these two decompositions, we get:

$$A + B = \sum_i \alpha_i R_i + \sum_j \beta_j S_j$$

   where $R_i \in \mathcal{A}_+(A+B)$ and $S_j \in \mathcal{A}_+(A+B)$ for all i and j. This decomposition shows that

$$\tau_+(A+B) \leq \tau_+(A) + \tau_+(B).$$

2. We now prove the property for τ_+^{sos}. Let (t, X) and (t', X') be the optimal points of the semidefinite program (3.17) for A and B respectively (i.e., $t = \tau_+^{\text{sos}}(A)$ and $t' = \tau_+^{\text{sos}}(B)$). It is not hard to see that $(t+t', X+X')$ is feasible for the semidefinite program that defines $\tau_+^{\text{sos}}(A+B)$ (in particular we use the fact that since A and B are nonnegative we have $A_{ij}^2 + B_{ij}^2 \leq (A_{ij} + B_{ij})^2$). Thus this shows that $\tau_+^{\text{sos}}(A+B) \leq \tau_+^{\text{sos}}(A) + \tau_+^{\text{sos}}(B)$.

Product

1. We first show the property for τ_+. We need to show that $\tau_+(AB) \leq \min(\tau_+(A), \tau_+(B))$. To see why note that if $R \in \mathcal{A}_+(A)$, then $RB \in \mathcal{A}_+(AB)$. Thus if we have $A = \sum_i \alpha_i R_i$ with $R_i \in \mathcal{A}_+(A)$ and $\sum_i \alpha_i = \tau_+(A)$, then we get $AB = \sum_i \alpha_i R_i B$ where each $R_i B \in \mathcal{A}_+(AB)$ and thus $\tau_+(AB) \leq \sum_i \alpha_i = \tau_+(A)$. The same reasoning shows that $\tau_+(AB) \leq \tau_+(B)$, and thus we get $\tau_+(AB) \leq \min(\tau_+(A), \tau_+(B))$.

2. We now prove the property for τ_+^{sos}, i.e., we show $\tau_+^{\text{sos}}(AB) \leq \min(\tau_+^{\text{sos}}(A), \tau_+^{\text{sos}}(B))$. We will show here that $\tau_+^{\text{sos}}(AB) \leq \tau_+^{\text{sos}}(A)$, and a similar reasoning can then be used to show $\tau_+^{\text{sos}}(AB) \leq \tau_+^{\text{sos}}(B)$.

   Let (t, X) be the optimal point of the semidefinite program (3.17) that defines $\tau_+^{\text{sos}}(A)$, i.e., $t = \tau_+^{\text{sos}}(A)$. We will show that the pair $(t, \widetilde{X})$ with

   $$\widetilde{X} = (B^T \otimes I_m) X (B \otimes I_m),$$

   is feasible for the semidefinite program that defines $\tau_+^{\text{sos}}(AB)$ and thus this will show that $\tau_+^{\text{sos}}(AB) \leq t = \tau_+^{\text{sos}}(A)$.

   Observe that we have $\operatorname{vec}(AB) = (B^T \otimes I_m) \operatorname{vec}(A)$ thus:

   $$\begin{bmatrix} t & \operatorname{vec}(AB)^T \\ \operatorname{vec}(AB) & \widetilde{X} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & B^T \otimes I \end{bmatrix} \begin{bmatrix} t & \operatorname{vec}(A)^T \\ \operatorname{vec}(A) & X \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & B \otimes I \end{bmatrix}$$

   and thus this shows that the matrix

   $$\begin{bmatrix} t & \operatorname{vec}(AB)^T \\ \operatorname{vec}(AB) & \widetilde{X} \end{bmatrix}$$

   is positive semidefinite.

   Using the definition of Kronecker product one can verify that the entries of $\widetilde{X}$ are given by:

   $$\widetilde{X}_{ij,kl} = \sum_{\alpha,\beta=1}^{m'} B_{\alpha j} B_{\beta l} X_{i\alpha,k\beta}.$$

   Using this formula we easily verify that $\widetilde{X}$ satisfies the rank-one equality constraints:

   $$\widetilde{X}_{ij,kl} = \widetilde{X}_{il,kj}$$

   since X itself satisfies the constraints.

   Finally it remains to show that $\widetilde{X}_{ij,ij} \leq (AB)_{ij}^2$. For this we need the following simple lemma:

   Lemma 2. *Let (t, X) be a feasible point for the semidefinite program (3.17). Then $X_{ij,kl} \leq A_{ij} A_{kl}$ for any i, j, k, l.*

Proof. Consider the 2×2 principal submatrix of X:

$$\begin{bmatrix} X_{ij,ij} & X_{ij,kl} \\ X_{kl,ij} & X_{kl,kl} \end{bmatrix}.$$

We know that $X_{ij,ij} \leq A_{ij}^2$ and $X_{kl,kl} \leq A_{kl}^2$. Furthermore since X is positive semidefinite we have $X_{ij,ij}X_{kl,kl} - X_{ij,kl}^2 \geq 0$. Thus we get that:

$$X_{ij,kl}^2 \leq X_{ij,ij}X_{kl,kl} \leq (A_{ij}A_{kl})^2.$$

Thus since $A_{ij}A_{kl} \geq 0$ we have $X_{ij,kl} \leq A_{ij}A_{kl}$. □

Using this lemma we get:

$$\widetilde{X}_{ij,ij} = \sum_{\alpha,\beta=1}^{m'} B_{\alpha j}B_{\beta j}X_{i\alpha,i\beta} \leq \sum_{\alpha,\beta=1}^{m'} B_{\alpha j}B_{\beta j}A_{i\alpha}A_{i\beta} = ((AB)_{ij})^2$$

which is what we want.

Monotonicity

Here we prove the monotonicity property of τ_+ and τ_+^{sos}. More precisely we show that if $A \in \mathbb{R}_+^{m\times n}$ is a nonnegative matrix, and B is a submatrix of A, then $\tau_+(B) \leq \tau_+(A)$ and $\tau_+^{\text{sos}}(B) \leq \tau_+^{\text{sos}}(A)$.

1. We prove the claim first for τ_+. Let $I \subseteq [m]$ and $J \subseteq [n]$ such that $B = A[I, J]$ (i.e., B is obtained from A by keeping only the rows in I and the columns in J). Let $X \in \operatorname{conv} \mathcal{A}_+(A)$ such that $A = \tau_+(A)X$. Define $Y = X[I, J]$ and note that $Y \in \operatorname{conv}(\mathcal{A}_+(B))$. Furthermore observe that we have $B = A[I, J] = \tau_+(A)X[I, J] = \tau_+(A)Y$. Hence, since $Y \in \operatorname{conv} \mathcal{A}_+(B)$, this shows, by definition of $\tau_+(B)$ that $\tau_+(B) \leq \tau_+(A)$.

2. We prove the claim now for the semidefinite programming relaxation τ_+^{sos}. As above, let $I \subseteq [m]$ and $J \subseteq [n]$ such that $B = A[I, J]$. Let (t, X) be the optimal point in (3.17) for the matrix A. It is easy to see that $(t, X[I, J])$ is feasible for the semidefinite program (3.17) for the matrix $B = A[I, J]$. Thus this shows that $\tau_+^{\text{sos}}(B) \leq \tau_+^{\text{sos}}(A)$.

Block-diagonal matrices

We now turn to prove the following: if $A \in \mathbb{R}_+^{m\times n}$ and $B \in \mathbb{R}_+^{m'\times n'}$ are two nonnegative matrices and $A \oplus B$ is the block-diagonal matrix:

$$A \oplus B = \begin{bmatrix} A & 0 \\ 0 & B \end{bmatrix},$$

then

$$\tau_+(A \oplus B) = \tau_+(A) + \tau_+(B) \quad \text{and} \quad \tau_+^{\text{sos}}(A \oplus B) = \tau_+^{\text{sos}}(A) + \tau_+^{\text{sos}}(B)$$

1. We first prove the claim for the quantity τ_+. Observe that the set $\mathcal{A}_+(A \oplus B)$ is equal to:

$$\mathcal{A}_+(A \oplus B) = \left\{ \begin{bmatrix} R & 0 \\ 0 & 0 \end{bmatrix} \ : \ R \in \mathcal{A}_+(A) \right\} \cup \left\{ \begin{bmatrix} 0 & 0 \\ 0 & R' \end{bmatrix} \ : \ R' \in \mathcal{A}_+(B) \right\}. \quad (3.33)$$

Indeed any element in $\mathcal{A}_+(A \oplus B)$ must have the off-diagonal blocks equal to zero (since the off-diagonal blocks of $A \oplus B$ are zero), and thus by the rank-one constraint at least one of the diagonal blocks is also equal to zero. Thus this shows that $\mathcal{A}_+(A \oplus B)$ decomposes as in (3.33).

We start by showing $\tau_+(A \oplus B) \geq \tau_+(A) + \tau_+(B)$. Let $Y \in \operatorname{conv} \mathcal{A}_+(A \oplus B)$ such that

$$A \oplus B = \tau_+(A \oplus B)Y.$$

Since $\mathcal{A}_+(A \oplus B)$ has the form (3.33), we know that Y can be decomposed as:

$$Y = \sum_{i=1}^{r} \lambda_i \begin{bmatrix} R_i & 0 \\ 0 & 0 \end{bmatrix} + \sum_{j=1}^{r'} \mu_j \begin{bmatrix} 0 & 0 \\ 0 & R'_{i'} \end{bmatrix},$$

where $R_i \in \mathcal{A}_+(A), R'_j \in \mathcal{A}_+(B)$ and $\sum_i \lambda_i + \sum_j \mu_j = 1$ with $\lambda, \mu \geq 0$. Note that since $A \oplus B = \tau_+(A \oplus B)Y$ we have:

$$A = \tau_+(A \oplus B) \sum_{i=1}^{r} \lambda_i R_i,$$

and

$$B = \tau_+(A \oplus B) \sum_{j=1}^{r'} \mu_j R'_j.$$

Hence $\tau_+(A) \leq \tau_+(A \oplus B) \sum_{i=1}^{r} \lambda_i$ and $\tau_+(B) \leq \tau_+(A \oplus B) \sum_{j=1}^{r'} \mu_j$ and we thus get:

$$\tau_+(A) + \tau_+(B) \leq \tau_+(A \oplus B) \left(\sum_{i=1}^{r} \lambda_i + \sum_{j=1}^{r'} \mu_j \right) = \tau_+(A \oplus B).$$

We now prove the converse inequality, i.e., $\tau_+(A \oplus B) \leq \tau_+(A) + \tau_+(B)$: Let $t = \tau_+(A)$, $t' = \tau_+(B)$ and $X \in \operatorname{conv} \mathcal{A}_+(A), X' \in \operatorname{conv} \mathcal{A}_+(B)$ such that $A = tX$ and $B = t'X'$. Define the matrix

$$Y = \begin{bmatrix} \frac{t}{t+t'} X & 0 \\ 0 & \frac{t'}{t+t'} X' \end{bmatrix},$$

and note that $A \oplus B = (t + t')Y$. If we show that $Y \in \operatorname{conv} \mathcal{A}_+(A \oplus B)$ then

this will show that $\tau_+(A \oplus B) \leq t + t'$. We can rewrite Y as:

$$Y = \frac{t}{t+t'} \begin{bmatrix} X & 0 \\ 0 & 0 \end{bmatrix} + \frac{t'}{t+t'} \begin{bmatrix} 0 & 0 \\ 0 & X' \end{bmatrix},$$

and it is easy to see from this expression that $Y \in \operatorname{conv} \mathcal{A}_+(A \oplus B)$.

We have thus proved that $\tau_+(A \oplus B) = \tau_+(A) + \tau_+(B)$.

2. We now prove the claim for the SDP relaxation τ_+^{sos}. Let $a = \operatorname{vec}(A)$ and $b = \operatorname{vec}(B)$. Since the matrix $A \oplus B$ has zeros on the off-diagonal, the SDP defining $\tau_+^{\text{sos}}(A \oplus B)$ can be simplified and we can eliminate the zero entries from the program. One can show that after the simplification we get that $\tau_+^{\text{sos}}(A \oplus B)$ is equal to the value of the SDP below:

$$\begin{aligned} \text{minimize} \quad & t \\ \text{subject to} \quad & \begin{bmatrix} t & a^T & b^T \\ a & X & 0 \\ b & 0 & X' \end{bmatrix} \succeq 0 \\ & X_{ij,ij} \leq A_{ij}^2 \quad \forall (i,j) \in [m] \times [n] \\ & X_{ij,kl} = X_{il,kj} \quad 1 \leq i < k \leq m \text{ and } 1 \leq j < l \leq n \\ & X'_{i'j',i'j'} \leq B_{i'j'}^2 \quad \forall (i',j') \in [m'] \times [n'] \\ & X'_{i'j',k'l'} = X_{i'l',k'j'} \quad 1 \leq i' < k' \leq m' \text{ and } 1 \leq j' < l' \leq n' \end{aligned} \tag{3.34}$$

It is well-known (see e.g., [56]) that the following equivalence always holds:

$$\begin{bmatrix} t & a^T & b^T \\ a & X & 0 \\ b & 0 & X' \end{bmatrix} \succeq 0 \iff \exists t_1, t_2 \,:\, t_1 + t_2 = t, \quad \begin{bmatrix} t_1 & a^T \\ a & X \end{bmatrix} \succeq 0, \quad \begin{bmatrix} t_2 & b^T \\ b & X' \end{bmatrix} \succeq 0$$

Using this equivalence, the semidefinite program (3.34) becomes:

$$\begin{aligned} \text{minimize} \quad & t_1 + t_2 \\ \text{subject to} \quad & \begin{bmatrix} t_1 & a^T \\ a & X \end{bmatrix} \succeq 0 \\ & X_{ij,ij} \leq A_{ij}^2 \quad \forall (i,j) \in [m] \times [n] \\ & X_{ij,kl} = X_{il,kj} \quad 1 \leq i < k \leq m \text{ and } 1 \leq j < l \leq n \\ & \begin{bmatrix} t_2 & b^T \\ b & X' \end{bmatrix} \succeq 0 \\ & X'_{i'j',i'j'} \leq B_{i'j'}^2 \quad \forall (i',j') \in [m'] \times [n'] \\ & X'_{i'j',k'l'} = X_{i'l',k'j'} \quad 1 \leq i' < k' \leq m' \text{ and } 1 \leq j' < l' \leq n' \end{aligned} \tag{3.35}$$

The semidefinite program is decoupled and it is easy to see that its value is equal to $\tau_+^{\text{sos}}(A) + \tau_+^{\text{sos}}(B)$.

3.5.2 Proof of Theorem 7 on the relation between τ_+ and $\tau_+^{\mathbf{sos}}$ with combinatorial bounds

Proof of Theorem 7. 1. We prove first that $\tau_+(A) \geq \chi_{\text{frac}}(\text{RG}(A))$. For convenience, we recall below the definitions of $\tau_+(A)$ and $\chi_{\text{frac}}(\text{RG}(A))$:

$\tau_+(A)$	$\chi_{\text{frac}}(\text{RG}(A))$
$\min \quad t$	$\min \quad t$
$\text{s.t.} \quad A \in t\,\text{conv}(\mathcal{A}_+(A))$	$\text{s.t.} \quad \exists Y \in t\,\text{conv}(\mathcal{A}_B(A))$
	$\text{s.t. } \forall (i,j), \;\; A_{i,j} > 0 \;\Rightarrow\; Y_{i,j} \geq 1$

Let $t = \tau_+(A)$ and $X \in \text{conv}(\mathcal{A}_+(A))$ such that $A = tX$. Consider the decomposition of X:

$$X = \sum_{k=1}^{r} \lambda_k X_k,$$

where $X_k \in \mathcal{A}_+(A)$, $\lambda_k \geq 0$ and $\sum_{k=1}^r \lambda_k = 1$. Let $R_k = \text{supp}(X_k)$ (i.e., R_k is obtained by replacing the nonzero entries of X_k with ones) and observe that $R_k \in \mathcal{A}_B(A)$. Define

$$Y = t \sum_{k=1}^{r} \lambda_k R_k \in t\,\text{conv}(\mathcal{A}_B(A))$$

Observe that for any (i,j) such that $A_{i,j} > 0$ we have:

$$Y_{i,j} = t \sum_{k : X_k[i,j] > 0} \lambda_k \underbrace{R_k[i,j]}_{=1} \overset{(a)}{\geq} t \sum_{k : X_k[i,j]} \lambda_k \frac{X_k[i,j]}{A_{i,j}} \overset{(b)}{=} \frac{A_{i,j}}{A_{i,j}} = 1$$

where in (a) we used the fact that $X_k \leq A$ (by definition of $X_k \in \mathcal{A}_+(A)$) and in (b) we used the fact that $A = t\sum_k \lambda_k X_k$. Thus this shows that (t, Y) is feasible for the optimization program defining $\chi_{\text{frac}}(\text{RG}(A))$ and thus we have $\chi_{\text{frac}}(\text{RG}(A)) \leq t = \tau_+(A)$.

2. We now show that $\tau_+^{\text{sos}}(A) \geq \overline{\vartheta}(\text{RG}(A))$. For convenience, we recall the two SDPs (3.19) and (3.22) that define $\tau_+^{\text{sos}}(A)$ and $\overline{\vartheta}(\text{RG}(A))$ below (note the constraint $X_{ij,ij} = A_{ij}^2$ in the SDP on the left appears as an inequality constraint in (3.19)—in fact it is not hard to see that with an equality constraint we get

the same optimal value):

$$\begin{aligned}
\tau_+^{\text{sos}}(A) = \quad & \min. \quad t \\
& \text{s.t.} \quad \begin{bmatrix} t & \pi(A)^T \\ \pi(A) & X \end{bmatrix} \succeq 0 \\
& \quad \forall (i,j) \text{ s.t. } A_{i,j} > 0 : \quad X_{ij,ij} = A_{ij}^2 \\
& \quad \forall (1,1) \le (i,j) < (k,l) \le (m,n) : \\
& \qquad \begin{cases} \text{if } A_{i,l}A_{k,j} = 0 \ : \ X_{ij,kl} = 0 & \text{(a)} \\ \text{if } A_{i,j}A_{k,l} = 0 \ : \ X_{il,kj} = 0 & \text{(b)} \\ \text{else} \quad X_{ij,kl} - X_{il,kj} = 0 & \text{(c)} \end{cases}
\end{aligned}$$

$$\begin{aligned}
\overline{\vartheta}(\text{RG}(A)) = \quad & \min. \quad t \\
& \text{s.t.} \quad \begin{bmatrix} t & 1^T \\ 1 & X \end{bmatrix} \succeq 0 \\
& \quad \forall (i,j) \text{ s.t. } A_{i,j} > 0 : \quad X_{ij,ij} = 1 \\
& \quad \forall (1,1) \le (i,j) < (k,l) \le (m,n) : \\
& \qquad \begin{cases} \text{if } A_{i,l}A_{k,j} = 0 \ : \ X_{ij,kl} = 0 & \text{(a')} \\ \text{if } A_{i,j}A_{k,l} = 0 \ : \ X_{il,kj} = 0 & \text{(b')} \end{cases}
\end{aligned}$$

Observe that the two semidefinite programs are very similar except that $\tau_+^{\text{sos}}(A)$ has more constraints than $\overline{\vartheta}(\text{RG}(A))$; cf. constraints (c) for $\tau_+^{\text{sos}}(A)$. To show that $\tau_+^{\text{sos}}(A) \ge \overline{\vartheta}(\text{RG}(A))$, let (t, X) be the solution of the SDP on the left for $\tau_+^{\text{sos}}(A)$. We will construct X' such that (t, X') is feasible for the SDP on the right and thus this will show that $\tau_+^{\text{sos}}(A) \ge \overline{\vartheta}(\text{RG}(A))$. Define X' by:

$$X' = \text{diag}(\pi(A))^{-1} X \, \text{diag}(\pi(A))^{-1}.$$

We show that (t, X') is feasible for the SDP on the right: Note that:

$$\begin{bmatrix} t & 1^T \\ 1 & X' \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & \text{diag}(\pi(A))^{-1} \end{bmatrix} \begin{bmatrix} t & \pi(A)^T \\ \pi(A) & X \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & \text{diag}(\pi(A))^{-1} \end{bmatrix} \succeq 0$$

Second we clearly have $X'_{ij,ij} = A_{ij}^{-2} X_{ij,ij} = 1$. Finally constraints (a') and (b') are also clearly true. Thus this shows that (t, X') is feasible for the SDP of $\overline{\vartheta}(\text{RG}(A))$ and thus $\overline{\vartheta}(\text{RG}(A)) \le t = \tau_+^{\text{sos}}(A)$.

□

Chapter 4

Equivariant semidefinite lifts

In this chapter we study the class of semidefinite lifts that respect symmetries, which we call "equivariant SDP lifts". We prove a structure theorem that gives a characterization of these lifts in terms of sum-of-squares certificates of the facet inequalities from an invariant subspace. We use this characterization to prove lower bounds on the size of equivariant SDP lifts for certain families of polytopes (cut polytope, parity polytope, regular polygons). This chapter is mostly based on the paper [37], except Section 4.6 on regular polygons which is based on part of [36].

4.1 Preliminaries: definitions and examples

Many polytopes $P \subset \mathbb{R}^n$ of interest in discrete and combinatorial optimization have symmetries, i.e., they are invariant under a certain group of transformations of $\mathbb{R}^n$. For such symmetric polytopes one may be interested in lifts that "respect" this symmetry. In the context of linear programming, such lifts were first studied by Yannakakis [101] where he showed that any symmetric LP lift of the matching polytope and of the traveling salesman polytope must have exponential size. In the more recent works [63, 47, 86, 50], it was shown that the symmetry requirement can have a significant impact on the size of lifts, i.e., there are polytopes (like the permutahedron for example) where there is a large gap between the smallest LP lift and the smallest *symmetric* LP lift. The recent work of Chan el al. [21] establishes, among others, a strong connection between symmetric LP lifts and the Sherali-Adams hierarchy: it is shown that the approximation quality of any polynomial-size symmetric LP for the maximum cut problem can be achieved by a constant number of rounds of the Sherali-Adams hierarchy[1].

Equivariant SDP lifts The works cited above studied the symmetry requirement in the context of LP lifts. Here we are interested in *SDP lifts* that respect the symmetries of the polytope P. Intuitively a SDP lift $P = \pi(Q)$ where $Q = \mathbf{S}^d_+ \cap L$ respects the symmetry of P if any transformation $g \in GL_n(\mathbb{R})$ which leaves P invariant can be lifted to a transformation $\Phi(g) \in GL(\mathbf{S}^d)$ that preserves both $\mathbf{S}^d_+$ and L, and so that the following *equivariance* relation holds: for any $y \in Q$, $\pi(\Phi(g)y) = g\pi(y)$. It is known that the transformations of $\mathbf{S}^d$ which leave the psd cone $\mathbf{S}^d_+$ invariant are precisely congruence transformations, see e.g. [99, Theorem 9.6.1]. This motivates the following definition of *equivariant SDP lift* which we adopt here:

Definition 9. Let $P \subset \mathbb{R}^n$ be a polytope invariant under the action of a group $G \subset GL_n(\mathbb{R})$. Assume $P = \pi(\mathbf{S}^d_+ \cap L)$ is a SDP lift of P of size d. The lift is called *G-equivariant* if there is a group homomorphism $\rho : G \to GL_d(\mathbb{R})$ such that the

[1] In fact the paper [21] establishes a connection between general LP formulations, possibly non-symmetric, and the Sherali-Adams hierarchy however the results are stronger in the case of symmetric LPs.

following two conditions hold:
(i) The subspace L is invariant under congruence by $\rho(g)$, for all $g \in G$:

$$\rho(g) Y \rho(g)^T \in L \quad \forall g \in G,\ \forall Y \in L. \tag{4.1}$$

(ii) The following equivariance relation holds:

$$\pi\left(\rho(g) Y \rho(g)^T\right) = g\pi(Y) \quad \forall g \in G,\ \forall Y \in \mathbf{S}^d_+ \cap L. \tag{4.2}$$

Observe that the notion of equivariant lift is defined with respect to a group G which leaves P invariant and this group does not have to be the full automorphism group of P. Indeed one may be interested in lifts that preserve only a certain subset of the symmetries of P, but not all of them. One example we discuss in detail later is the *parity polytope* which is invariant under permutation of coordinates as well as under certain sign switches. In Section 4.4 we mention two examples of well-known lifts of the parity polytope which are equivariant with respect to one set of transformations but not the other.

Examples To illustrate the definition of equivariant SDP lift, we now give a simple example of an equivariant SDP lift and another example of a SDP lift that does not satisfy the definition of equivariance.

Example 7. **An equivariant SDP lift of the square** $[-1,1]^2$
Consider the SDP lift of the square $[-1,1]^2$ that we saw in Example 2, Chapter 2. We recall the lift here for convenience:

$$[-1,1]^2 = \left\{ (x_1, x_2) \in \mathbb{R}^2 \ :\ \exists u \in \mathbb{R} \ \begin{bmatrix} 1 & x_1 & x_2 \\ x_1 & 1 & u \\ x_2 & u & 1 \end{bmatrix} \succeq 0 \right\}. \tag{4.3}$$

We can write this lift in standard form as $[-1,1]^2 = \pi(\mathbf{S}^3_+ \cap L)$ where L and π are given by:

$$L = \{X \in \mathbf{S}^3 \ :\ X_{11} = X_{22} = X_{33} = 1\} \quad \text{and} \quad \pi(X) = (X_{12}, X_{13}) \in \mathbb{R}^2.$$

The symmetry group of the square $[-1,1]^2$ is the dihedral group of order 8, denoted D_8. To show that the lift (4.3) is D_8-equivariant, consider the group homomorphism $\rho : D_8 \to GL_3(\mathbb{R})$ defined by:

$$\rho(g) = \begin{bmatrix} 1 & 0 \\ 0 & g \end{bmatrix} \quad \forall g \in D_8.$$

It is easy to see that the congruence operation by $\rho(g)$ stabilizes the subspace L, and that the following equivariance relation holds:

$$\pi(\rho(g) X \rho(g)^T) = g\pi(X) \quad \forall g \in D_8, \forall X \in \mathbf{S}^3_+ \cap L.$$

Thus this shows that the SDP lift (4.3) is D_8-equivariant. $\lozenge$

We now show an example of a non-equivariant SDP lift.

Example 8. **A nonequivariant SDP lift of the hyperboloid**
Let H be the hyperboloid in $\mathbb{R}^3$ defined by:

$$H = \{(x_1, x_2, x_3) \in \mathbb{R}^3 \ : \ x_1, x_2, x_3 \geq 0 \text{ and } x_1 x_2 x_3 \geq 1\}.$$

One can construct a SDP lift of H of size 6 as follows (see e.g., [10, page 261]):

$$\begin{aligned} H &= \left\{(x_1, x_2, x_3) \in \mathbb{R}^3 \ : \ \exists y, z \geq 0 \quad x_1 x_2 \geq y^2, \ x_3 \geq z^2, \ yz \geq 1\right\} \\ &= \left\{(x_1, x_2, x_3) \in \mathbb{R}^3 \ : \ \exists y, z \qquad \begin{bmatrix} x_1 & y \\ y & x_2 \end{bmatrix} \succeq 0, \ \begin{bmatrix} x_3 & z \\ z & 1 \end{bmatrix} \succeq 0, \ \begin{bmatrix} y & 1 \\ 1 & z \end{bmatrix} \succeq 0 \right\}. \end{aligned} \tag{4.4}$$

The hyperboloid H is clearly invariant under permutation of coordinates, i.e., for any permutation $\sigma \in \mathfrak{S}_3$ we have $(x_1, x_2, x_3) \in H \Rightarrow \sigma \cdot (x_1, x_2, x_3) \in H$. However the lift we just constructed does not respect this symmetry: indeed to construct the lift we imposed a particular ordering of the variables where the last coordinate x_3 does not play the same role as the first two coordinates x_1 and x_2. It is not difficult to formally show that the lift (4.4) does not satisfy Definition 9 of equivariance when $G = \mathfrak{S}_3$. Note however that the lift is equivariant with respect to permuting the coordinates x_1 and x_2. ◊

Relation with symmetric LP lift Symmetric LP lifts can be interpreted as equivariant SDP lifts, where each $\rho(g)$ consists of a permutation matrix. In fact, recall that an LP lift of a polytope P takes the form $P = \pi(\mathbb{R}^d_+ \cap L)$ where L is an affine subspace of $\mathbb{R}^d$ and $\pi : \mathbb{R}^d \to \mathbb{R}^n$ is a linear map. An LP lift $P = \pi(\mathbb{R}^d_+ \cap L)$ is called G-symmetric (or G-equivariant) if there exists $\theta : G \to \mathfrak{S}_d$ (where $\mathfrak{S}_d$ is the group of permutations on d elements) such that for any $y \in \mathbb{R}^d_+ \cap L$, $\pi(\theta(g) \cdot y) = g \cdot \pi(y)$. By working with diagonal matrices, any symmetric LP lift can be rewritten as an equivariant SDP lift: indeed, if $P = \pi(\mathbb{R}^d_+ \cap L)$ is a symmetric LP lift of P, then $P = \widetilde{\pi}(\mathbf{S}^d_+ \cap \widetilde{L})$ is an equivariant SDP lift where $\widetilde{L} = \{Y \in \mathbf{S}^d, Y \text{ is diagonal and } \operatorname{diag}(Y) \in L\}$ and $\widetilde{\pi} = \pi \circ \operatorname{diag}$ where $\operatorname{diag} : \mathbf{S}^d \to \mathbb{R}^d$ is the operator extracting the diagonal of a symmetric matrix. This SDP lift clearly satisfies the definition of equivariance where $\rho(g)$ is the permutation matrix associated to $\theta(g)$.

Related work In independent work, Lee et al. [77] have also considered symmetric SDP lifts. In their work however they adopted a definition of symmetric SDP lift that is different from the one we consider here. Since the focus of [77] is on constraint satisfaction problems, the authors only consider symmetry with respect to the permutation group. If P is a polytope invariant under permutation of coordinates then an SDP lift $P = \pi(\mathbf{S}^d_+ \cap L)$ is called symmetric in [77] if for any permutation σ of $\{1, \ldots, n\}$ there exists a $d \times d$ *permutation matrix* $\rho(\sigma)$ such that $\pi(\rho(\sigma) Y \rho(\sigma)^T) = \sigma \cdot \pi(Y)$ for all $Y \in \mathbf{S}^d_+ \cap L$. Note that this definition is more restrictive than ours since it requires $\rho(\sigma)$ to be a permutation matrix whereas in our Definition 9, we allow $\rho(\sigma)$ to be any invertible matrix in $GL_d(\mathbb{R})$. In this regard our framework is more general

and applies to a wider class of SDP lifts.

Organization of the chapter The chapter is organized as follows. In Section 4.2 we review some background material and terminology in representation theory that we use later in the chapter. In Section 4.3 we state and prove the main theorem (Theorem 10) which gives a characterization of equivariant SDP lifts in terms of sum-of-squares certificates of facet inequalities from an invariant subspace. We then apply our main theorem to three different examples of polytopes: the parity polytope (Section 4.4), the cut polytope (Section 4.5), and regular polygons (Section 4.6). For the parity polytope and the cut polytope we show that any equivariant SDP lift must have exponential size. For regular N-gons we show that any equivariant SDP lift must have size at least $\Omega(\log N)$.

4.2 Background: invariant subspaces and irreducible subspaces

We recall some basic facts concerning representation theory of finite groups which will be used later. We refer to [95] for a reference. Given a finite group G, a real finite-dimensional representation of G is a pair (V, ρ) where V is a real finite-dimensional vector space and $\rho : G \to GL(V)$ is a group homomorphism. Two representations (V_1, ρ_1) and (V_2, ρ_2) are called G-isomorphic if there is an isomorphism $f : V_1 \to V_2$ such that $f(\rho_1(g)x) = \rho_2(g)f(x)$ for all $x \in V_1$ and $g \in G$. A subspace W of V is an *invariant subspace* for the representation ρ if for any $x \in W$ and $g \in G$ we have $\rho(g)x \in W$. The representation (V, ρ) of G is called *irreducible* if it does not contain any invariant subspace except $\{0\}$ and V itself. Irreducible representations of a group G are the building blocks of any representation of G. The following result is a standard fact in representation theory: any finite-dimensional real representation (V, ρ) of G can be decomposed as

$$V = V_1 \oplus \cdots \oplus V_k \tag{4.5}$$

where each V_i is isomorphic to a direct sum of m_i copies of an irreducible representation W_i of G. This decomposition (4.5) is a canonical decomposition and is called the *isotypic decomposition* of V. It satisfies the following important property: if W is an irreducible subspace of V that is G-isomorphic to W_i then W is contained in V_i. The subspace V_i is called the isotypic component of the irreducible representation W_i in V.

This decomposition result can be used to prove the following proposition which will be needed later:

Proposition 3. *Let (V, ρ) be a real finite-dimensional representation of a finite group G and assume*

$$V = W_1 \oplus \cdots \oplus W_h$$

is a decomposition of V into irreducibles, i.e., each W_i is an irreducible subspace of V. Assume furthermore that the W_i are sorted in nondecreasing order of dimension, i.e., $\dim W_1 \leq \dim W_2 \leq \cdots \leq \dim W_h$. *Assume W is an invariant subspace of V with* $\dim W < \dim W_{i_0}$ *for some* $i_0 \in \{1, \ldots, h\}$. *Then necessarily W is contained in the direct sum* $W_1 \oplus \cdots \oplus W_{i_0-1}$.

Proof. Any irreducible subrepresentation of W is isomorphic to one of the W_i for some $i < i_0$. Thus W is contained in the direct sum of isotypic components of the W_i's for $i < i_0$, thus W is contained in $\bigoplus_{i=1}^{i_0-1} W_i$. □

The following well-known proposition will also be useful later.

Proposition 4. *Let $\rho : G \to GL_n(\mathbb{R})$ be a real finite-dimensional representation of a finite group G. Then there exists an invertible matrix Q such that $Q\rho(g)Q^{-1}$ is orthogonal for all $g \in G$.*

Proof. With the choice $Q = (\sum_{g \in G} \rho(g)\rho(g)^T)^{-1/2}$, one can easily verify that $Q\rho(g)Q^{-1}$ is orthogonal for all $g \in G$. □

4.3 Structure theorem

4.3.1 Sums of squares from invariant subspaces

We saw in Chapter 2 (see Theorem 5) that finding an SDP lift of $\text{conv}(X)$ can be reduced to finding sum of squares certificates of the facet inequalities of $\text{conv}(X)$ from a low-dimensional subspace V of $\mathbb{R}^X$. It is not difficult to turn this theorem into a theorem that produces G-equivariant SDP lifts, where G is a symmetry group for X. In fact one way to obtain such an equivariant lift is to require that the subspace V be *invariant* under the action of G. Here, the action of G on $\mathbb{R}^X$ is defined in the natural way as:

$$(g \cdot f)(x) = f(g^{-1} \cdot x)$$

for any $g \in G, f \in \mathbb{R}^X, x \in X$. A subspace V of $\mathbb{R}^X$ is called *G-invariant* if $g \cdot f \in V$ for any $f \in V$ and $g \in G$.

The following example illustrates the notion of invariant subspace.

Example 9. Consider the set $X \subset \mathbb{R}^2$ of N roots of unity, i.e.,

$$X = \{x_k = (\cos(2\pi k)/N, \sin(2\pi k/N)) : k \in \mathbb{Z}_N\}.$$

The symmetry group of X is the dihedral group which consists of N rotations and N reflections. The rotation of angle $2\pi t/N$, where $t \in \mathbb{Z}_N$ maps the point x_k to the point x_{k+t}. The reflection indexed by $t \in \mathbb{Z}_N$ maps point x_k to x_{t-k}. Since the elements of X are indexed by $\mathbb{Z}_N$ we can think of a function $f \in \mathbb{R}^X$ as a vector $f \in \mathbb{R}^{\mathbb{Z}_N}$. The action of a rotation of angle $2\pi t/N$ on f corresponds to *shifting* the components of f. This depicted in Figure 4-1.

Consider the subspace V_j of $\mathbb{R}^X \cong \mathbb{R}^{\mathbb{Z}_N}$ spanned by the two functions $c_j(k) = \cos(2j\pi k/N)$ and $s_j(k) = \sin(2j\pi k/N)$. We claim that V_j is invariant under the

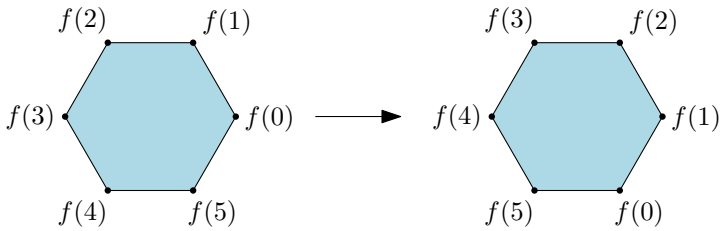

Figure 4-1: Action of rotation by $2\pi/N$ on a function $f \in \mathbb{R}^X$.

action of the dihedral group. Indeed if g is the rotation that maps x_k to x_{k+t} then we have:

$$\begin{aligned}(g \cdot c_j)(k) &= \cos(2j\pi(-t+k)/N) \\ &= \cos(2j\pi t/N)\cos(2j\pi k/N) + \sin(2j\pi t/N)\sin(2j\pi k/N) \\ &= \cos(2j\pi t/N)c_j(k) + \sin(2j\pi t/N)s_j(k).\end{aligned}$$

In other words we have $g \cdot c_j \in \operatorname{span}(c_j, s_j) = V_j$. Similarly we can show that if g is a reflection then $g \cdot c_j \in V_j$ and the same for s_j instead of c_j. Thus this shows that V_j is *invariant* with respect to the dihedral group of order $2N$. We will revisit this example of regular N-gons in more detail in Section 4.6 of this chapter. ◊

The following definition will be useful in the rest of this section.

Definition 10. Given $f \in \mathbb{R}^X$ and V a subspace of $\mathbb{R}^X$ we say that f is V-sos if there exist functions $h_1, \ldots, h_J \in V$ such that $f = \sum_{j=1}^J h_j^2$.

The following theorem should be compared to Theorem 5 (sufficiency part). It shows that if the subspace V is invariant under the action of the symmetry group G then the resulting SDP lift we obtain is G-equivariant.

Theorem 9. *Let $P \subset \mathbb{R}^n$ be a full-dimensional polytope with vertex set X and assume that X is invariant under the action of a group G. Assume furthermore that there is a G-invariant subspace V of $\mathbb{R}^X$ such that for any facet inequality $\ell \leq \ell_{\max}$ of P the function $\ell_{\max} - \ell|_X$ is V-sos on X. Then $\operatorname{conv}(X)$ has a G-equivariant SDP lift of size $\dim V$.*

Proof. See Section 4.8.1. □

Remark 5. Recall from Chapter 2 (see Section 2.3.2) that the well-known Lasserre/theta-body lift of $\operatorname{conv}(X)$ is obtained by taking V to be the space of polynomials of degree at most k on X (for a well-chosen k). Since the subspace of polynomials of degree at most k is G-invariant (for any linear action of a group G on $\mathbb{R}^n$) it follows that the Lasserre/theta-body lifts are G-equivariant.

4.3.2 Statement of structure theorem

A natural question is to ask whether Theorem 9 has a converse, namely whether for any small equivariant SDP lift there is an associated small invariant subspace of

functions from which we can certify nonnegativity of the facet inequalities using sums of squares. Our main theorem stated below shows that this is indeed the case.

Orbitopes We will focus on a family of polytopes that are symmetric "by construction" and are known as *orbitopes* [5, 4, 93]. These are constructed as follows: let G be a finite subgroup of $GL_n(\mathbb{R})$ and let $x_0 \in \mathbb{R}^n$. Define $X = G \cdot x_0 := \{g \cdot x_0 \ : \ g \in G\}$ to be the *orbit* of x_0 under G, and consider the associated *orbitope* defined as the convex hull of X:

$$P = \text{conv}(X) = \text{conv}(G \cdot x_0). \tag{4.6}$$

Orbitopes are symmetric by construction. For example they are clearly invariant under the action of G. One example of orbitopes are regular polygons in the plane: if we let $x_0 = (1, 0)$ and $G \cong \mathbb{Z}_N$ be the group of rotations of angle $\{2\pi k/N, k = 0, \ldots, N-1\}$ then $G \cdot x_0$ are the N roots of unity and $\text{conv}(G \cdot x_0)$ is the regular N-gon.

The next theorem gives a converse to Theorem 9. It shows that any equivariant SDP lift of an orbitope must be of the "sum-of-squares form" where the subspace V is G-invariant.

Theorem 10 (Structure theorem for equivariant SDP lifts)**.** *Let G be a finite group acting on $\mathbb{R}^n$ and let $X = G \cdot x_0$ where $x_0 \in \mathbb{R}^n$. Let $P = \text{conv}(X)$ and assume that P has a G-equivariant SDP lift of size d. Then there exists a G-invariant subspace V of $\mathbb{R}^X$ with $\dim V \leq d^3$ such that the following holds: for any facet inequality $\ell \leq \ell_{\max}$ of P the function $\ell_{\max} - \ell|_X$ is V-sos on X.*

Before presenting the proof of the theorem, we first give some comments:

- Theorem 10 is the analogue of Theorem 5 (necessity part) for the case of equivariant SDP lifts. It shows that if the SDP lift is equivariant, then the subspace V can be chosen to be G-invariant. The bound on the dimension is slightly worse (d^3 instead of d^2). In Section 4.3.3 we describe a specific situation where the bound can be improved.

- Theorem 10 shows that in order to obtain lower bounds on the sizes of equivariant SDP lifts, one has to study the structure of G-invariant subspaces of $\mathbb{R}^X$. For example, if one can show that such subspaces correspond to low-degree polynomials, then any lower bound on the Lasserre/theta-body hierarchy will yield a lower bound on G-equivariant SDP lifts. This will be our strategy to obtain lower bounds on equivariant SDP lifts of the parity polytope and cut polytope. In some cases however (e.g., regular polygons) low-dimensional invariant subspaces however do not necessarily correspond to low-degree polynomials and in this case proving lower bounds can be more challenging (cf. Section 4.6).

We conclude this section by presenting the proof of Theorem 10. The proof relies on a certain factorization theorem from [50] which we recall here and prove later for completeness.

Theorem 11 ([50, Theorem 2]). *Let G be a finite group acting on $\mathbb{R}^n$ and let $X = G \cdot x_0$ where $x_0 \in \mathbb{R}^n$. Assume $\operatorname{conv}(X) = \pi(\mathbf{S}^d_+ \cap L)$ is a G-equivariant SDP lift of $\operatorname{conv}(X)$ of size d, i.e., there is a homomorphism $\rho : G \to GL_d(\mathbb{R})$ such that conditions (i) and (ii) of Definition 9 hold. Then there exists a map $A : X \to \mathbf{S}^d_+$ with the following properties:*

(i) For any linear form ℓ on $\mathbb{R}^n$ there exists $B(\ell) \in \mathbf{S}^d_+$ such that if we let $\ell_{\max} := \max_{x \in X} \ell(x)$ we have:

$$\ell_{\max} - \ell(x) = \langle A(x), B(\ell) \rangle \quad \forall x \in X.$$

(ii) The map A satisfies the following equivariance relation:

$$A(g \cdot x) = \rho(g) A(x) \rho(g)^T \quad \forall x \in X,\ \forall g \in G.$$

In particular if H denotes the stabilizer of x_0, then we have:

$$A(x_0) = \rho(h) A(x_0) \rho(h)^T \quad \forall h \in H. \tag{4.7}$$

Furthermore, the representation $\rho : G \to GL_d(\mathbb{R})$ can be taken to be orthogonal, i.e., $\rho(g) \in O_d(\mathbb{R})$ for all $g \in G$.

Proof. The proof is given in Section 4.8.2. □

We now proceed to the proof of the structure theorem, Theorem 10.

Proof of Theorem 10. Assume we have a G-equivariant SDP lift of size d of $\operatorname{conv}(X)$. By the factorization theorem (Theorem 11), we know that there exist maps $A : X \to \mathbf{S}^d_+$ and $B : (\mathbb{R}^n)^* \to \mathbf{S}^d_+$ such that for any linear form $\ell \in (\mathbb{R}^n)^*$ we have:

$$\ell_{\max} - \ell(x) = \langle A(x), B(\ell) \rangle \quad \forall x \in X \tag{4.8}$$

where $\ell_{\max} := \max_{x \in X} \ell(x)$. Furthermore, since the lift is G-equivariant, the map A satisfies the equivariance relation

$$A(g \cdot x) = \rho(g) A(x) \rho(g)^T \quad \forall x \in X \ \ \forall g \in G \tag{4.9}$$

where $\rho : G \to O(\mathbb{R}^d)$ is a group homomorphism.

Let H be the stabilizer of x_0, i.e., $H = \{g \in G : g \cdot x_0 = x_0\}$. Note that the set X can be identified with G/H, the set of left cosets of H. Furthermore the left action of G on X is isomorphic to the left action of G on G/H. For simplicity of notation, we will thus think of functions on X as functions on G/H, or equivalently, as functions on G that are constant on the left cosets of H. For example since the point x_0 is identified with the left coset $1_G H$ of H, we will write $A(1_G)$ instead of $A(x_0)$.

We first show how to construct the subspace V of $\mathbb{R}^X \cong \mathbb{R}^{G/H}$ and then we prove that it satisfies the properties of the statement.

- *Definition of V*: Let

$$A(1_G) = \sum_{i=1}^{r} \lambda_i P_{W_i}$$

be an eigendecomposition of $A(1_G)$, where each P_{W_i} is an orthogonal projection on the eigenspace W_i. Observe that from Equation (4.7) we have $A(1_G) = \rho(h)A(1_G)\rho(h)^T$ for all $h \in H$. Since $\rho(h)$ is orthogonal, this means that $A(1_G)$ commutes with $\rho(h)$, and so $\rho(h)W_i = W_i$ for each eigenspace W_i of $A(1_G)$, which also implies that $\rho(h)P_{W_i}\rho(h)^T = P_{W_i}$. An important consequence of this is that the functions $g \mapsto \rho(g)P_{W_i}\rho(g)^T$ are constant on the left cosets of H, thus we can think of them as functions on G/H. Let V be the subspace of $\mathbb{R}^{G/H}$ spanned by the matrix entries of $x \in G/H \mapsto \rho(x)P_{W_i}\rho(x)^T$, namely

$$V = \operatorname{span}\Big\{ x \in G/H \mapsto (\rho(x)P_{W_i}\rho(x)^T)_{k,l},\ i = 1, \dots, r \ \text{ and } \ k, l = 1, \dots, d \Big\}. \tag{4.10}$$

- It is clear that V is G-invariant (since ρ is a homomorphism) and that $\dim V \leq d^3$. It thus remains to show that $\ell_{\max} - \ell$ is V-sos on X for any $\ell \in (\mathbb{R}^n)^*$. We know from (4.8) that there exists $B = B(\ell) \in \mathbf{S}^d_+$ such that $\ell_{\max} - \ell(x) = \langle A(x), B \rangle$ for any $x \in X$. Now for any $x \in G/H$ we have:

$$\langle A(x), B \rangle = \langle \rho(x)A(1_G)\rho(x)^T, B \rangle = \sum_{i=1}^{r} \lambda_i \langle \rho(x)P_{W_i}\rho(x)^T, B \rangle. \tag{4.11}$$

Note that even though $\rho(x)$ is *not* well-defined when $x \in G/H$, the quantities $\rho(x)A(1_G)\rho(x)^T$ and $\rho(x)P_{W_i}\rho(x)^T$ are well-defined and do not depend on the representative of the coset $x \in G/H$ (cf. previous remarks). Observe that since P_{W_i} is an orthogonal projection and $\rho(x)$ is orthogonal, we have $(\rho(x)P_{W_i}\rho(x)^T)^2 = \rho(x)P_{W_i}\rho(x)^T$. Thus continuing from Equation (4.11) we can write:

$$\langle A(x), B \rangle = \sum_{i=1}^{r} \lambda_i \langle (\rho(x)P_{W_i}\rho(x)^T)^2, B \rangle = \sum_{i=1}^{r} \lambda_i \| L^T \rho(x) P_{W_i} \rho(x)^T \|_F^2$$

where L is such that $B = L^T L$. Since each entry function of

$$x \in G/H \mapsto L^T \rho(x) P_{W_i} \rho(x)^T$$

lives in V, this shows that $\langle A(x), B \rangle$ is a sum-of-squares of functions in V, which is what we wanted. □

4.3.3 Groups with a product structure

In the case where the group G has a certain product structure, one can strengthen the conclusion of Theorem 10 with a better bound on the dimension of V. We thus assume in this section that the group G has the following property:

Assumption 1 (Product structure)**.** *Any element $g \in G$ can be written in a unique way as $g = nh$ where $n \in N, h \in H$ where N and H are subgroups of G that satisfy the following:*

- *N is a normal subgroup of G, i.e., for any $n \in N$ we have $gng^{-1} \in N$ for all $g \in G$*
- *H stabilizes x_0, i.e., $h \cdot x_0 = x_0$ for all $h \in H$*

In group-theoretic terms the conditions above can be summarized by saying that G is the *semidirect product* of N and H, i.e., $G = N \rtimes H$ and that H stabilizes x_0. The following example shows that the symmetry group of the cube has such a product structure.

Example 10 (The hypercube). Let $x_0 = (1, \ldots, 1) \in \mathbb{R}^n$ and let $G \subset GL_n(\mathbb{R})$ be the group of signed permutations, i.e., the group of permutation matrices where nonzero entries are either $+1$ or -1. The orbit $G \cdot x_0$ is $\{-1, 1\}^n$. Note that G has the product structure described above where

$$N = \{\mathrm{diag}(\epsilon) : \epsilon \in \{-1, 1\}^n\}$$
$$H = \text{permutation matrices.}$$

Indeed any signed permutation matrix can be written as a product nh where $n \in N$ and $h \in H$. Furthermore H stabilizes x_0 since $x_0 = (1, \ldots, 1)$. Finally N is normal because if h is any permutation matrix and $n \in N$ then $hnh^{-1} \in N$. ◊

We now state our structure theorem when the group G has a product structure.

Theorem 12 (Structure theorem; special case with product structure)**.** *Let G be a finite group acting on $\mathbb{R}^n$ and let $X = G \cdot x_0$ where $x_0 \in \mathbb{R}^n$. Assume that G satisfies Assumption 1: namely $G = N \rtimes H$ where N is a normal subgroup of G and H stabilizes x_0. Assume P has a G-equivariant SDP lift of size d. Then there exists a G-invariant subspace V of $\mathbb{R}^X$ with the following properties:*

(i) For any facet inequality $\ell \leq \ell_{\max}$ of $P = \mathrm{conv}(X)$ the nonnegative function $\ell_{\max} - \ell|_X$ is V-sos on X.

(ii) $\dim V \leq \alpha_N(d) \cdot d$ where $\alpha_N(d)$ is the largest dimension of any real irreducible representation of N of dimension $\leq d$ (in particular $\dim V \leq d^2$).

Remark 6. Note that the worst case bound on the dimension of V here is d^2 whereas in Theorem 10 it is d^3. Furthermore in the examples that we will consider later in the chapter the quantity $\alpha_N(d)$ will be equal to 1 in which case the bound reduces simply to d.

Proof of Theorem 12. Assume we have a G-equivariant SDP lift of size d of $\mathrm{conv}(X)$. By the factorization theorem (Theorem 11), we know that there exist maps $A : X \to \mathbf{S}^d_+$ and $B : (\mathbb{R}^n)^* \to \mathbf{S}^d_+$ such that for any linear form $\ell \in (\mathbb{R}^n)^*$ we have:

$$\ell_{\max} - \ell(x) = \langle A(x), B(\ell) \rangle \quad \forall x \in X \tag{4.12}$$

where $\ell_{\max} := \max_{x \in X} \ell(x)$. Furthermore, since the lift is G-equivariant, the map A satisfies the equivariance relation

$$A(g \cdot x) = \rho(g) A(x) \rho(g)^T \quad \forall x \in X \ \ \forall g \in G \tag{4.13}$$

where $\rho : G \to GL(\mathbb{R}^d)$ is a group homomorphism (in fact we can choose ρ to take values in $O(\mathbb{R}^d)$ however we will not need this in this proof).

Note that since G has the product structure $G = N \rtimes H$ and H is the stabilizer of x_0 we can identify X with N, and the left action of G on X is isomorphic to the left action of G on N defined by $g \cdot x = nhxh^{-1} \in N$, where $g = nh \in G$ and $x \in N$. Thus in the rest of the proof we will think of functions on X as functions on N. For example since the point x_0 is identified with 1_N, we will write $A(1_N)$ instead of $A(x_0)$.

We now define the subspace V of $\mathbb{R}^X \cong \mathbb{R}^N$ and then we show it has the required properties.

• *Definition of* V: Let V be the subspace of $\mathbb{R}^N$ spanned by the matrix entry functions of $\rho|_N$, i.e.,

$$V = \operatorname{span}\Big\{ x \in N \mapsto \rho(x)_{ij}, i, j = 1, \ldots, d \Big\}.$$

• We need to show that V is a G-invariant subspace and that Properties (i) and (ii) in the statement of the theorem are satisfied.

∗ To see why V is G-invariant, note that for any $x \in N$ and $g = nh \in G$ we have $\rho(g \cdot x) = \rho(nhxh^{-1}) = \rho(nh)\rho(x)\rho(h^{-1})$ thus for any i, j the function $x \mapsto \rho(g \cdot x)_{ij}$ is a linear combination of the functions $x \mapsto \rho(x)_{k,l}$. This shows that V is G-invariant.

∗ To prove that $\dim V \leq \alpha_N(d) \cdot d$, observe that if $n_1, \ldots, n_k$ are the dimensions of the irreducible components of the representation $\rho|_N : N \to GL(\mathbb{R}^d)$ of N, then the matrices $\rho(x)$ $(x \in N)$ are all, up to a global change of basis, block-diagonal with blocks of size $n_1, \ldots, n_k$. Thus we have $\dim V \leq \sum_i n_i^2 \leq \sum_i n_i \alpha_N(d) = \alpha_N(d) \cdot d$ since each $n_i \leq \alpha_N(d)$ and $\sum_i n_i = d$.

∗ Finally, it remains to prove Property (i). Let $\ell \in (\mathbb{R}^n)^*$ and let $B = B(\ell) \in \mathbf{S}^d_+$ such that (4.12) is true. Note that for any $x \in N$ we have:

$$\langle A(x), B \rangle = \langle \rho(x) A(1_N) \rho(x)^T, B \rangle = \mathbf{Tr}(\rho(x) A(1_N) \rho(x)^T B). \tag{4.14}$$

Since $A(1_N)$ and B are positive semidefinite matrices, we can write $A(1_N) = L_A L_A^T$ and $B = L_B L_B^T$ for some matrices L_A and L_B. Then we have:

$$\langle A(x), B \rangle = \mathbf{Tr}(\rho(x) L_A L_A^T \rho(x)^T L_B L_B^T) = \| L_B^T \rho(x) L_A \|_F^2.$$

Since each entry function of $x \mapsto L_B^T \rho(x) L_A$ lives in V, it follows that $x \mapsto \langle A(x), B \rangle$ has a sum-of-squares representation with functions from V. This completes the proof.

□

Remark 7. Observe that the theorem above could be stated more generally without using the language of *lifts*. Assume $p(x)$ is a function that is nonnegative on X and has an *equivariant certificate of nonnegativity* of the form

$$p(x) = \langle A(x), B \rangle \quad \forall x \in X$$

where $B \in \mathbf{S}^d_+$ and $A : X \to \mathbf{S}^d_+$ satisfies the equivariance relation:

$$A(g \cdot x) = \rho(g) A(x) \rho(g)^T \quad \forall x \in X, \ \forall g \in G.$$

Then the arguments in the proofs above show that $p(x)$ must be a sum of squares of functions from a G-invariant subspace V of $\mathbb{R}^X$ whose dimension is bounded by a certain function of d (d^3 in the setting of Theorem 10 and $\alpha_N(d)d$ in the setting of Theorem 12). In the theorems above, the function $p(x)$ corresponds to the facet-defining linear function $p(x) = \ell_{\max} - \ell(x)$ but the proofs did not use this specific form of the function $p(x)$. This will be useful later when analyzing approximate SDP lifts of the cut polytope (Section 4.5).

4.3.4 Illustration: the square $[-1, 1]^2$

In this section we illustrate how one can use the structure theorems to derive a lower bound on equivariant SDP lifts of the square $P = [-1, 1]^2$. We will show, via Theorem 12 that the square $[-1, 1]^2$ does not admit a G-equivariant SDP lift of size smaller than 3 where $G = D_8$ is the symmetry group of the square (the dihedral group of order 8).

First recall that in Example 7 we gave a G-equivariant SDP lift of $P = [-1, 1]^2$ of size 3. We can apply Theorem 12 to this lift (recall from Example 10 that the symmetry group here has the required product structure): Theorem 12 says that there must exist a G-invariant subspace V of $\mathbb{R}^{\{-1,1\}^2}$ with the following properties:

(i) Any facet inequality $\ell(x) \leq \ell_{\max}$ of P has a sum-of-squares certificate with functions from V:

$$\ell_{\max} - \ell(x) = \sum_j f_j(x)^2 \quad \forall x \in \{-1, 1\}^2$$

where $f_j \in V$.

(ii) $\dim V \leq 1 \cdot 3 = 3$ (indeed $\alpha_N(3) = 1$ since N is isomorphic to $\mathbb{Z}_2^2$ for which all the real irreducible representations have dimension one).

It is actually not difficult to construct this subspace V explicitly. In fact, we have already constructed it in Example 4 from Chapter 2: the subspace V of polynomials of degree at most 1 on $\{-1, 1\}^2$ (i.e., $V = \text{span}(1, x_1, x_2)$) satisfies point (i) above, has dimension 3 and is clearly G-invariant.

Now one may wonder if there exists an equivariant SDP lift of the square $P = [-1,1]^2$ of size 2. Using Theorem 12 this would mean that there exists a G-invariant subspace V of $\mathbb{R}^{\{-1,1\}^2}$ of dimension ≤ 2 which allows us to certify the four facet inequalities of P using sum-of-squares. Later in the chapter we will study in more detail the space of functions on the hypercube $\{-1,1\}^n$ and their invariant subspaces, cf. Lemma 4. Using this lemma one can actually show that such a subspace V of dimension ≤ 2 cannot exist, ruling out the existence of G-equivariant SDP lifts of size 2 of the square $[-1,1]^2$.

Remark 8. Actually it is known that there does not exist any SDP lift (even a nonequivariant one) of the square $[-1,1]^2$ of size 2. Indeed it was shown in [51] that any SDP lift of a full-dimensional polytope P in $\mathbb{R}^n$ must have size at least $n+1$.

4.4 Application 1: the parity polytope

In this section we derive lower bounds on the size of equivariant SDP lifts of the parity polytope using the structure theorem.

4.4.1 Definitions

Define EVEN_n to be the set of points $x \in \{-1,1\}^n$ that have an even number of -1's, i.e.:

$$\text{EVEN}_n = \left\{ x \in \{-1,1\}^n \ : \ \prod_{i=1}^n x_i = 1 \right\}. \tag{4.15}$$

The convex hull of EVEN_n is called the *parity polytope* and is denoted PAR_n:

$$\text{PAR}_n = \text{conv}(\text{EVEN}_n).$$

Symmetry group The group of transformations that leave the parity polytope invariant consist of *evenly signed permutation matrices*, i.e., permutation matrices where each entry can be either $+1$ or -1 with the constraint that the total number of -1's is even. We denote by G_{parity} this group. We prove in this section an exponential lower bound on the size of G_{parity}-equivariant SDP lifts of the parity polytope.

Theorem 13. *Any G_{parity}-equivariant SDP lift of PAR_n for $n \geq 8$ must have size $\geq \binom{n}{\lceil n/4 \rceil}$.*

Product structure of the symmetry group Note that the symmetry group of the parity polytope has the product structure corresponding to Section 4.3.3. The reasoning is very similar to the hypercube group considered in Example 10. In fact let $N_{\text{parity}} \subset GL_n(\mathbb{R})$ be the subgroup of $GL_n(\mathbb{R})$ consisting of diagonal matrices with $+1$ or -1 on the diagonal and with an even number of -1's, i.e.:

$$N_{\text{parity}} = \{\text{diag}(x) \ : \ x \in \text{EVEN}_n\}.$$

It is not difficult to see that any element of $g \in G_{\text{parity}}$ can be written in a unique way as $g = \epsilon h$ where $\epsilon \in N_{\text{parity}}$ and h is a permutation matrix. Since N_{parity} is normal in G and permutation matrices stabilize $x_0 = (1, \ldots, 1)$ it follows that Assumption 1 is satisfied.

Facet description of the parity polytope When $n > 2$, the parity polytope is a full-dimensional polytope in $\mathbb{R}^n$. It has the following description using linear inequalities (see [101, 60]):

$$\text{PAR}_n = \Big\{ x \in \mathbb{R}^n \ : \ -1 \leq x \leq 1, \ \sum_{i \in A^c} x_i - \sum_{i \in A} x_i \leq n - 2 \quad \forall A \subseteq [n], |A| \text{ odd} \Big\}. \tag{4.16}$$

If $n \geq 4$ each of the $2n + 2^{n-1}$ inequalities are facet-defining. If $n = 3$ the inequalities $-1 \leq x \leq 1$ are redundant giving the simpler description with 4 facets

$$\begin{aligned} \text{PAR}_3 = \{ x \in \mathbb{R}^n \ : \quad x_1 + x_2 - x_3 &\leq 1 \\ x_1 - x_2 + x_3 &\leq 1 \\ -x_1 + x_2 + x_3 &\leq 1 \\ -x_1 - x_2 - x_3 &\leq 1 \}. \end{aligned} \tag{4.17}$$

Nonequivariant polynomial-size lifts of the parity polytope Polynomial-size lifts of the parity polytope have been known since the original paper of Yannakakis [101]. In fact there are two known LP lifts of the parity polytope of size respectively $O(n^2)$ and $O(n)$. The two lifts respect some of the symmetry of the parity polytope however none of them is equivariant with respect to the *full* symmetry group $G_{\text{parity}} = N_{\text{parity}} \rtimes \mathfrak{S}_n$.

- The lift of size $O(n^2)$ given by Yannakakis [101] relies on the observation that

$$\text{PAR}_n = \text{conv}\left(\bigcup_{k \text{ even}} S_k \right)$$

  where S_k are the "slices" of the hypercube defined by the equation $1^T x = n - 2k$:

$$\begin{aligned} S_k &= \text{conv}\{ x \in \{-1, 1\}^n \ : \ x \text{ has exactly } k \text{ components equal to } -1 \} \\ &= \{ x \in [-1, 1]^n \ : \ 1^T x = n - 2k \}. \end{aligned}$$

  Since each S_k has a simple description using only $O(n)$ linear inequalities, this can be used to construct a lift of PAR_n of size $O(n^2)$. One can easily verify that this lift is equivariant with respect to permutation of the coordinates. One can show however that this lift is *not* equivariant with respect to switching an even number of signs. Intuitively, the reason behind this is that the operation of switching signs does not preserve the slices S_k.

- There is a smaller yet less well known LP lift of the parity polytope due to [20, Section 2.6.3] (see also [62]) which has size $O(n)$. This lift is equivariant with

respect to switching an even number of signs, however it is *not* equivariant with respect to the permutation action. The key observation behind this LP lift is that $(x_1, x_2, \ldots, x_n) \in \mathrm{EVEN}_n$ if and only if there exists $z \in \{-1, 1\}$ such that $(x_1, x_2, z) \in \mathrm{EVEN}_3$ and $(z, x_3, \ldots, x_n) \in \mathrm{EVEN}_{n-1}$ (simply take $z = x_1 x_2$). In fact one can establish an analog of this statement for the parity polytope:

$$\begin{aligned} \mathrm{PAR}_n = \{x \in \mathbb{R}^n : \exists z \text{ s.t. } & (x_1, x_2, z) \in \mathrm{PAR}_3 \\ & \text{and } (z, x_3, \ldots, x_n) \in \mathrm{PAR}_{n-1}\}. \end{aligned} \tag{4.18}$$

Repeatedly applying (4.18) shows that

$$\begin{aligned} \mathrm{PAR}_n = \Big\{ x \in \mathbb{R}^n : \exists z_2, z_3, \ldots, z_{n-2} \text{ s.t.} \\ & (x_1, x_2, z_2) \in \mathrm{PAR}_3,\ (z_{n-2}, x_{n-1}, x_n) \in \mathrm{PAR}_3, \\ & \text{and } (z_i, x_{i+1}, z_{i+1}) \in \mathrm{PAR}_3 \quad \text{for } i \in \{2, 3, \ldots, n-3\} \Big\}. \end{aligned}$$

This description shows that PAR_n is the projection of a polytope with $4(n-2)$ facets. It is not too hard to show that this lift is actually equivariant with respect to switching an even number of signs. However one can see that this lift is not equivariant with respect to permutations since we have broken permutation symmetry by imposing a particular ordering on the variables.

4.4.2 Invariant subspaces of functions on EVEN_n

In order to understand equivariant SDP lifts of PAR_n, we need to understand G_{parity}-invariant subspaces of $\mathbb{R}^{\mathrm{EVEN}_n}$. This is the object of this section.

If $I \subseteq [n]$ define the monomial map $\mathbb{R}^n \ni x \mapsto x^I := \prod_{i \in I} x_i$. We can regard these as functions on EVEN_n by simply restricting their domain. When we do so, we write them as e_I so that:

$$e_I : \mathrm{EVEN}_n \to \mathbb{R}, \quad e_I(x) = x^I.$$

When working on the hypercube $\{-1, 1\}^n$ it is well-known that the functions $(e_I)_{I \subseteq [n]}$ form a basis of functions on the hypercube. When working on EVEN_n we have the additional fact that $e_I = e_{I^c}$ and so one only needs half of these functions. This is made precise in the next theorem.

Proposition 5. *Let $n \geq 1$.*

- *If n is odd, then the functions e_I for $|I| < n/2$ form a basis of $\mathbb{R}^{EVEN_n}$.*
- *If n is even, then the functions e_I with $|I| < n/2$ together with the functions $(e_I + e_{I^c})/2$ for $|I| = n/2$ constitute a basis of $\mathbb{R}^{EVEN_n}$.*

Proof. Given $a \in \mathrm{EVEN}_n$, let $1_a \in \mathbb{R}^{\mathrm{EVEN}_n}$ be the indicator function for a, i.e., $1_a(x) = 1$ if $x = a$ and $1_a(x) = 0$ otherwise. Clearly the family of functions

$\{1_a\}_{a\in \mathrm{EVEN}_n}$ forms a basis of $\mathbb{R}^{\mathrm{EVEN}_n}$. Observe that 1_a can be written as the following polynomial:

$$1_a(x) = \frac{1}{2^n}(1 + a_1x_1)\dots(1 + a_nx_n) \quad \forall x \in \mathrm{EVEN}_n.$$

Furthermore, the following identities are true on EVEN_n: $x_i^2 = 1$ and $x^I = x^{I^c}$ for any $x \in \mathrm{EVEN}_n$. If we expand the polynomial expression for 1_a above using the previous identities we see that, when n is odd, 1_a is a linear combination of the square-free monomials x^I for $|I| < n/2$. When n is even any monomial of the form x^I where $|I| = n/2$ can be rewritten as $(x^I + x^{I^c})/2$. Thus this shows that the functions given in the statement of the proposition form a generating set for $\mathbb{R}^{\mathrm{EVEN}_n}$. Since the number of such functions is $2^{n-1} = \dim \mathbb{R}^{\mathrm{EVEN}_n}$, they form a basis of $\mathbb{R}^{\mathrm{EVEN}_n}$. □

Given $0 \le k < n/2$, let $\mathrm{Pol}_k(\mathrm{EVEN}_n)$ be the subspace of $\mathbb{R}^{\mathrm{EVEN}_n}$ of homogeneous polynomials of degree k, i.e.,

$$\mathrm{Pol}_k(\mathrm{EVEN}_n) = \mathrm{span}\{e_I : |I| = k\}.$$

If $k = n/2$ let $\mathrm{Pol}_k(\mathrm{EVEN}_n)$ be the subspace of $\mathbb{R}^{\mathrm{EVEN}_n}$ spanned by the $(e_I + e_{I^c})/2$ with $|I| = n/2$. So we have:

$$\dim \mathrm{Pol}_k(\mathrm{EVEN}_n) = \begin{cases} \binom{n}{k} & \text{if } k < n/2 \\ \frac{1}{2}\binom{n}{n/2} & \text{if } k = n/2. \end{cases}$$

Proposition 5 shows that the space $\mathbb{R}^{\mathrm{EVEN}_n}$ decomposes as:

$$\mathbb{R}^{\mathrm{EVEN}_n} = \mathrm{Pol}_0(\mathrm{EVEN}_n) \oplus \dots \oplus \mathrm{Pol}_{\lfloor n/2 \rfloor}(\mathrm{EVEN}_n).$$

For future reference we let $\mathrm{Pol}_{\le k}(\mathrm{EVEN}_n)$ be the space of polynomials on EVEN_n of degree at most k:

$$\mathrm{Pol}_{\le k}(\mathrm{EVEN}_n) = \bigoplus_{i=0}^{k} \mathrm{Pol}_i(\mathrm{EVEN}_n).$$

Irreducible subspaces We are interested in subspaces of $\mathbb{R}^{\mathrm{EVEN}_n}$ that are G_{parity}-invariant. Recall that G_{parity} is the group of evenly signed permutations. Thus a subspace V of $\mathbb{R}^{\mathrm{EVEN}_n}$ is G_{parity}-invariant if for any $f \in V$, and any $\epsilon \in \{-1, +1\}^n$ such that $\prod_{i=1}^n \epsilon_i = 1$, and any $\sigma \in \mathfrak{S}_n$ the function:

$$x \mapsto f(\epsilon_1 x_{\sigma(1)}, \dots, \epsilon_n x_{\sigma(n)})$$

is also in V. Recall that an invariant subspace V is called *irreducible* if it does not contain any nontrivial invariant subspace, i.e., if W is an invariant subspace of V, then $W = \{0\}$ or $W = V$. It is clear that the subspaces $\mathrm{Pol}_k(\mathrm{EVEN}_n)$ are G_{parity}-invariant. The next result shows that these subspaces are actually irreducible. For the statement to follow we use the following notation (where n and k are two integers

such that $k \leq n/2$):

$$D_{n,k} := \begin{cases} \binom{n}{k} & \text{if } n \text{ is odd} \\ \min\left(\binom{n}{k}, \frac{1}{2}\binom{n}{n/2}\right) & \text{if } n \text{ is even.} \end{cases} \tag{4.19}$$

Lemma 3. *Under the action of G_{parity}, $\mathbb{R}^{EVEN_n}$ decomposes into irreducible invariant subspaces as*

$$\mathbb{R}^{EVEN_n} = \mathrm{Pol}_0(EVEN_n) \oplus \cdots \oplus \mathrm{Pol}_{\lfloor n/2 \rfloor}(EVEN_n).$$

Hence if V is a G_{parity}-invariant subspace of $\mathbb{R}^{EVEN_n}$ with $\dim(V) < D_{n,k}$ then $V \subseteq \mathrm{Pol}_{\leq k-1}(EVEN_n)$.

Proof. See Section 4.8.3. □

4.4.3 Lower bound on equivariant SDP lifts

In this section we prove our main theorem concerning equivariant SDP lifts of the parity polytope. Lemma 3 tells us that low-dimensional invariant subspaces of $\mathbb{R}^{\mathrm{EVEN}_n}$ correspond to low-degree polynomials. This allows us to show that if the parity polytope admits a small equivariant SDP lift, then a few levels of the Lasserre/theta-body hierarchy are enough to be exact. This is the object of the next theorem.

Theorem 14. *Assume PAR_n has a G_{parity} equivariant SDP lift of size d where d satisfies $d < D_{n,k}$ for some $k \leq n/2$. Then the $(k-1)$'st Lasserre/theta-body relaxation is exact, i.e., $\mathrm{TH}_{k-1}(EVEN_n) = PAR_n$.*

Proof. This is a direct consequence of the structure theorem and of Lemma 3. Assume we have a G_{parity}-equivariant SDP lift of PAR_n of size d. We can apply Theorem 12 with $P = \mathrm{PAR}_n$ and $G = G_{\mathrm{parity}} = N_{\mathrm{parity}} \rtimes \mathfrak{S}_n$. Since $N_{\mathrm{parity}} \cong \mathbb{Z}_2^{n-1}$, all the real irreducible representations of N_{parity} are one-dimensional. Thus Theorem 12 says that there exists a G_{parity}-invariant subspace V of $\mathbb{R}^{\mathrm{EVEN}_n}$ with $\dim V \leq d$ such that for any linear form $\ell \in (\mathbb{R}^n)^*$ we have that

$$\ell_{\max} - \ell \text{ is } V\text{-sos on } \mathrm{EVEN}_n \tag{4.20}$$

where $\ell_{\max} := \max_{x \in \mathrm{EVEN}_n} \ell(x)$. In Lemma 3 we showed that such an invariant subspace, when $d < D_{n,k}$, is composed entirely of polynomials of degree at most $k-1$, i.e. V is a subspace of $\mathrm{Pol}_{\leq k-1}(\mathrm{EVEN}_n)$. Thus this shows that $\mathrm{TH}_{k-1}(\mathrm{EVEN}_n) = \mathrm{PAR}_n$. □

Theta-rank To obtain a lower bound on equivariant SDP lifts of the parity polytope we need a lower bound on its *theta-rank*, i.e., the number of levels required by the Lasserre/theta-body hierarchy for exactness (cf. Section 2.3.2). Such a lower bound was already obtained in [48, Corollary 5.7] where it was shown that the theta-rank of the parity polytope is exactly $\lceil n/4 \rceil$.

Proposition 6 (Collorary 5.7 in [48]). *The theta-rank of the parity polytope PAR_n is exactly $\lceil n/4 \rceil$.*

Proof. We include a self-contained proof for the lower bound in Section 4.8.4 for completeness. □

Exponential lower bounds We are now ready to prove Theorem 13 giving an exponential lower bound on the size of equivariant SDP lifts of the parity polytope.

Theorem 13. *Any G_{parity}-equivariant SDP lift of PAR_n for $n \geq 8$ must have size $\geq \binom{n}{\lceil n/4 \rceil}$.*

Proof. We apply Theorem 14 with $k = \lceil n/4 \rceil$. By Proposition 6, we know that the theta-body relaxation of order $\lceil n/4 \rceil - 1$ is *not* exact. Thus this means that any G_{parity}-equivariant SDP lift of PAR_n must have size $d \geq D_{n,\lceil n/4 \rceil}$. One can then easily check from the definition of $D_{n,k}$ that when $n \geq 8$ we have $D_{n,\lceil n/4 \rceil} \geq \binom{n}{\lceil n/4 \rceil}$. □

Approximate lifts Note that, using Remark 7 from the previous section, one can actually state a more general theorem relating *approximate* equivariant SDP lifts and the sum-of-squares hierarchy. Indeed one can prove the following:

Theorem 15. *Assume $\widehat{P} \subseteq \mathbb{R}^n$ is an outer-approximation of PAR_n (i.e., $PAR_n \subseteq \widehat{P}$) and assume $\widehat{P}$ has a G_{parity}-equivariant SDP lift of size d. If $d < D_{n,k}$ for some $k \leq n/2$ then necessarily*

$$PAR_n \subseteq \text{TH}_{k-1}(EVEN_n) \subseteq \widehat{P}$$

where $\text{TH}_{k-1}(EVEN_n)$ is the $(k-1)$'st Lasserre/theta-body relaxation for the parity polytope.

Proof. The proof is very similar to the proof of Theorem 14 above. Given a linear form $\ell \in (\mathbb{R}^n)^*$ let $\ell_{\max}$ and $\widehat{\ell_{\max}}$ be respectively the maximum of ℓ on PAR_n and $\widehat{P}$. Note that $\ell_{\max} \leq \widehat{\ell_{\max}}$ since $\text{PAR}_n \subseteq \widehat{P}$, and hence the linear function $\widehat{\ell_{\max}} - \ell(x)$ is nonnegative on EVEN_n. Since $\widehat{P}$ has a G_{parity}-equivariant SDP lift of size d, and since $\text{EVEN}_n \subseteq \widehat{P}$ one can show (using a simple generalization of Theorem 11) that we have an equivariant certificate of nonnegativity of $\widehat{\ell_{\max}} - \ell$ on EVEN_n of the form:

$$\widehat{\ell_{\max}} - \ell(x) = \langle A(x), B \rangle \quad \forall x \in \text{EVEN}_n$$

where A satisfies the equivariance relation $A(g \cdot x) = \rho(g) A(x) \rho(g)^T$ for all $x \in \text{EVEN}_n$, $g \in G_{\text{parity}}$. Thus by Remark 7, we know that $\widehat{\ell_{\max}} - \ell(x)$ is a sum-of-squares of functions in a G_{parity}-invariant subspace V of dimension $\leq d$. Thus using Lemma 3 below it holds that $\widehat{\ell_{\max}} - \ell(x)$ is a sum-of-squares of functions of degree $\leq k-1$ on EVEN_n.

This is true for any facet-defining linear form ℓ of $\widehat{P}$ thus, by the definition of the theta-body relaxation (cf. Equation (2.21)) we have $\text{TH}_{k-1}(\text{EVEN}_n) \subseteq \widehat{P}$. □

4.5 Application 2: the cut polytope

In this section we use the structure theorem to derive lower bounds on equivariant SDP lifts of the cut polytope.

4.5.1 Definitions and symmetry group

The maximum cut problem on a graph $G = (V, E)$ with $V = \{1, \ldots, n\}$ and weights w_{ij} for $ij \in E$ is the problem of labeling each vertex $i \in V$ with a label $x_i = +1$ or $x_i = -1$ in such a way that the total weight of edges connecting two vertices with a different label is maximized. This problem can be written as follows:

$$\begin{array}{ll} \text{maximize} & \sum_{ij \in E} w_{ij}(1 - x_i x_j)/2 \\ \text{subject to} & x \in \{-1, 1\}^n. \end{array} \tag{4.21}$$

Note that for a given labeling $x_i = \pm 1$ of vertices, the quantity $(1 - x_i x_j)/2$ is equal to 1 if i and j have different labels, and 0 otherwise. The formulation (4.21) shows that the maximum cut problem is the problem of maximizing a *quadratic* form on the hypercube $\{-1, 1\}^n$. Using standard techniques (e.g., as outlined in Chapter 1), one can convert this problem into a *linear* program, by working in a lifted space. Indeed it is not hard to see that the problem (4.21) is equivalent to the problem below:

$$\begin{array}{ll} \text{maximize} & \sum_{ij \in E} w_{ij}(1 - X_{ij})/2 \\ \text{subject to} & X = xx^T \quad \text{for some } x \in \{-1, 1\}^n. \end{array} \tag{4.22}$$

Note that the objective is now linear in the variable X. Define the cut polytope CUT_n as the convex hull of all outer products xx^T for $x \in \{-1, 1\}^n$:

$$\mathrm{CUT}_n = \mathrm{conv}\left\{xx^T \ : \ x \in \{-1, 1\}^n\right\}. \tag{4.23}$$

The formulation (4.22) shows that the maximum cut problem is a linear program over the cut polytope CUT_n. Note that the cut polytope is a $n(n-1)/2$-dimensional polytope in the space $\mathbf{S}^n$ of $n \times n$ symmetric matrices.

Symmetries of the hypercube and the cut polytope Let

$$C_n = \{-1, 1\}^n$$

be the vertices of the hypercube in $\mathbb{R}^n$. The symmetry group of C_n consists of *signed permutation matrices*, i.e., permutation matrices where each nonzero entry is ± 1. Note each signed permutation matrix g can be written in a unique way as $g = \epsilon h$ where ϵ is a ± 1-diagonal matrix and h is a permutation matrix. In fact we saw in Example 10 that this symmetry group has the product structure described in Section 4.3.3.

The group G_{cube} acts on the space of $n \times n$ symmetric matrices by congruence

transformations as follows:

$$g \cdot X := gXg^T \quad \forall g \in G_{\text{cube}},\ \forall X \in \mathbf{S}^n. \tag{4.24}$$

It is easy to verify that CUT_n is invariant under this action of G_{cube}.

In this section we prove that any G_{cube}-equivariant SDP lift of CUT_n must have exponential size.

Theorem 16. *Any G_{cube}-equivariant SDP lift of CUT_n must have size $\geq \binom{n}{\lfloor n/4 \rfloor}$.*

4.5.2 Sum-of-squares relaxations

In this section we review the sum-of-squares relaxations of the cut polytope as described for example in [72]. The construction we describe here actually applies to general polytopes P of the form:

$$P = \text{conv}\left\{xx^T \ :\ x \in X\right\} \subset \mathbf{S}^n \tag{4.25}$$

where X is a finite set in $\mathbb{R}^n$. The construction of the relaxation is very similar to the one described in Section 2.2.5 of Chapter 2, except that we work with quadratic forms instead of linear forms. Note that a polytope P of the form (4.25) can be seen as the set of second-order moments of probability distributions on X:

$$P = \left\{ (E(e_{ij}))_{ij} \ : E \in (\mathbb{R}^X)^* \text{ where } \right.$$
$$\left. E(f) = \int_X f(x) d\mu(x) \text{ for some prob. measure } \mu \text{ on } X \right\} \tag{4.26}$$

where e_{ij} is the element of $\mathbb{R}^X$ defined by $e_{ij}(x) = x_i x_j$. Given a subspace V of $\mathbb{R}^X$ we can thus define the following relaxation of P in the same way we did in Section 2.2.5:

$$\text{TH}_V^{(2)}(X) := \left\{ (E(e_{ij}))_{ij} \ :\ E \in (\mathbb{R}^X)^* \text{ where } E(1) = 1,\ E(f^2) \geq 0\ \forall f \in V \right\}. \tag{4.27}$$

By comparing (4.26) and (4.27) it is clear that $P \subseteq \text{TH}_V^{(2)}(X)$. One can show, like in Theorem 5 that the relaxation is tight $\text{TH}_V^{(2)}(X) = P$ if for any quadratic form q on $\mathbb{R}^n$, $q_{\max} - q$ is V-sos, where $q_{\max} = \max_{x \in X} q(x)$.

When $X = \{-1, 1\}^n$, the convex set $\text{TH}_V^{(2)}(\{-1,1\}^n)$ is a relaxation of the cut polytope. When $V = \text{Pol}_{\leq k}(\{-1,1\}^n)$ is the space of polynomials of degree at most k, we will denote the relaxation by $Q_k(\text{CUT}_n)$:

$$Q_k(\text{CUT}_n) = \text{TH}_V^{(2)}(\{-1,1\}^n) \quad \text{where} \quad V = \text{Pol}_{\leq k}(\{-1,1\}^n). \tag{4.28}$$

The relaxation $Q_k(\text{CUT}_n)$ is known as the "node-based" relaxation of the cut polytope

and is usually described in the literature in terms of explicit moment matrices, see e.g., [73]. In fact if we use the basis of $\mathrm{Pol}_{\leq k}(\{-1,1\}^n)$ formed by square-free monomials of degree up to k, we get that $Q_k(\mathrm{CUT}_n)$ can be written as:

$$Q_k(\mathrm{CUT}_n) = \left\{ z \in \mathbb{R}^{\binom{n}{2}} \ : \ \exists (y_I)_{|I|\leq 2k} \text{ such that } \begin{array}{l} y_\emptyset = 1 \\ y_{ij} = z_{ij}, \ \ \forall i < j \\ \mathcal{M}_k(y) \succeq 0 \end{array} \right\} \tag{4.29}$$

where $(y_I)_{|I|\leq 2k}$ is a vector indexed by subsets $I \subseteq [n]$ of cardinality $\leq 2k$ and where $\mathcal{M}_k(y)$ is the moment matrix associated to y defined by:

$$\mathcal{M}_k(y)_{I,J} = y_{I \triangle J} \quad \forall I, J \subseteq [n], \ |I|, |J| \leq k$$

where $I \triangle J$ denotes symmetric difference. Laurent showed in [72] that when $k \leq \lfloor n/2 \rfloor$ we have $Q_k(\mathrm{CUT}_n) \neq \mathrm{CUT}_n$.

Theorem 17 (Laurent, [72])**.** *For $k \leq \lceil n/2 \rceil - 1$, the inclusion $CUT_n \subset Q_k(CUT_n)$ is strict.*

Laurent conjectured in [72] that the relaxation is actually tight for $k = \lceil n/2 \rceil$. We will prove this conjecture in the next chapter (see Theorem 26).

4.5.3 Invariant subspaces of functions on the hypercube

In order to understand equivariant SDP lifts of CUT_n, we need to understand G_{cube}-invariant subspaces of $\mathbb{R}^{C_n}$ where $C_n = \{-1,1\}^n$. It is known that any function $f \in \mathbb{R}^{C_n}$ can be seen as a square-free polynomial of degree at most n. Let $\mathrm{Pol}_k(C_n)$ be the space of homogeneous square-free polynomials of degree k. Note that $\dim \mathrm{Pol}_k(C_n) = \binom{n}{k}$ and that:

$$\mathbb{R}^{C_n} = \mathrm{Pol}_0(C_n) \oplus \cdots \oplus \mathrm{Pol}_n(C_n).$$

We are interested in subspaces of $\mathbb{R}^{C_n}$ that are G_{cube}-invariant. Recall that G_{cube} is the group of signed permutation matrices. Thus a subspace V of $\mathbb{R}^{C_n}$ is G_{cube}-invariant if for any $f \in V, \epsilon \in \{-1,+1\}^n, \sigma \in \mathfrak{S}_n$ the function:

$$x \mapsto f(\epsilon_1 x_{\sigma(1)}, \ldots, \epsilon_n x_{\sigma(n)})$$

is also in V.

It is clear that the subspaces $\mathrm{Pol}_k(C_n)$ are G_{cube}-invariant. The next result shows that these subspaces are actually irreducible under the action of G_{cube}.

Lemma 4. *Under the action of G_{cube}, $\mathbb{R}^{C_n}$ decomposes into irreducible invariant subspaces as*

$$\mathbb{R}^{C_n} = \mathrm{Pol}_0(C_n) \oplus \cdots \oplus \mathrm{Pol}_n(C_n).$$

Furthermore, suppose $k < n/2$. Then $\mathrm{Pol}_{n-k}(C_n) \cong e_n(x)\,\mathrm{Pol}_k(C_n)$ where $e_n(x) = x_1 \cdots x_n$ is the n'th elementary symmetric polynomial. Hence if V is a G_{cube}-invariant

subspace with $\dim(V) < \binom{n}{k}$ *then every* $f \in V$ *has the form*

$$f(x) = g(x) + e_n(x)h(x) \tag{4.30}$$

where $g(x)$ *and* $h(x)$ *have degree* $\leq k-1$.

Proof. See Section 4.8.5. □

4.5.4 Lower bound on equivariant SDP lifts

In this section we prove our main theorem concerning equivariant SDP lifts of the cut polytope.

We begin by a theorem relating equivariant SDP lifts of CUT_n and the sum-of-squares hierarchy of CUT_n. Lemma 4 tells us that low-dimensional invariant subspaces of $\mathbb{R}^{C_n}$ correspond, essentially, to low-degree polynomials. This allows us to show that if the cut polytope admits a small equivariant SDP lift, then a few levels of the Lasserre/theta-body hierarchy for the cut polytope are enough to be exact. The next result makes this claim formal. Note that its proof requires an additional argument (compared to the proof of Theorem 14 for the parity polytope) in order to take care of the term $e_n(x)h(x)$ in Equation (4.30).

Theorem 18. *Assume* CUT_n *has a* G_{cube}*-equivariant SDP lift of size* d *where* $d < \binom{n}{k}$ *for some* $k \leq n/2$*. Then the* $(k-1)$*'st sum-of-squares relaxation of* $CUT_{\lfloor n/2 \rfloor}$ *is exact, i.e.,* $Q_{k-1}(CUT_{\lfloor n/2 \rfloor}) = CUT_{\lfloor n/2 \rfloor}$.

Proof. Assume we have a G_{cube}-equivariant SDP lift of CUT_n of size d. Using arguments very similar to Theorem 12 (where linear forms are replaced by quadratic forms) we can show that there exists a G_{cube}-invariant subspace V of $\mathbb{R}^{C_n}$ with $\dim V \leq d$ such that for any quadratic form q on n variables with $q_{\max} := \max_{x \in C_n} q(x)$ we have:

$$q_{\max} - q(x) = \sum_i f_i(x)^2 \quad \forall x \in C_n \tag{4.31}$$

where each $f_i \in V$. In Lemma 4 (cf. below) we show that such an invariant subspace of dimension $d < \binom{n}{k}$, is composed entirely of polynomials of the form $g(x)+e_n(x)h(x)$ where g and h are polynomials of degree at most $k-1$ and $e_n(x) = x_1 \cdots x_n$ is the n'th elementary symmetric polynomial.

We can use this to show that the $(k-1)$'st sos relaxation of $\mathrm{CUT}_{\lfloor n/2 \rfloor}$ is exact. Assume for simplicity that n is even, $n = 2m$ (the argument for n odd is very similar). Let q be an arbitrary quadratic form on m variables. We will show that $q_{\max} - q(x)$ is a sum-of-squares of polynomials of degree at most k on C_m (i.e., it is $\mathrm{Pol}_{\leq k}(C_m)$-sos). Define $\widehat{q}$ a quadratic form in $n = 2m$ variables by:

$$\widehat{q}(x_1, \ldots, x_n) = q(x_1, \ldots, x_m).$$

Note that the polynomial $\widehat{q}$ does not depend on $x_{m+1}, \ldots, x_n$ and note also that $\max_{x \in C_n} \widehat{q}(x) = \max_{x \in C_m} q(x)$, i.e., $\widehat{q}_{\max} = q_{\max}$. From Equation (4.31) we know that

$\widehat{q}_{\max} - \widehat{q}(x)$ admits a sum-of-squares decomposition where each sos term lives in the subspace V, i.e., we have:

$$\widehat{q}_{\max} - \widehat{q}(x) = \sum_j (\widehat{g}_j(x) + e_n(x)\widehat{h}_j(x))^2 \quad \forall x \in C_n \tag{4.32}$$

where $\widehat{g}_j \in \mathrm{Pol}_{\leq k-1}(C_n)$ and $\widehat{h}_j \in \mathrm{Pol}_{\leq k-1}(C_n)$ and $e_n(x)$ is the n'th elementary symmetric polynomial $e_n(x) = x_1 \cdots x_n$. If we plug $x_{m+1} = x_1, x_{m+2} = x_2, \ldots, x_{2m} = x_m$ in Equation (4.32) we get:

$$q_{\max} - q(x) = \sum_j (g_j(x) + h_j(x))^2 \quad \forall x \in C_m \tag{4.33}$$

where we used the fact that $e_n(x_1, \ldots, x_m, x_1, \ldots, x_m) = 1$ for all $x \in C_m$ and where we let

$$\begin{aligned} g_j(x_1, \ldots, x_m) &= \widehat{g}_j(x_1, \ldots, x_m, x_1, \ldots, x_m) \\ h_j(x_1, \ldots, x_m) &= \widehat{h}_j(x_1, \ldots, x_m, x_1, \ldots, x_m) \end{aligned} \quad \forall (x_1, \ldots, x_m) \in C_m.$$

Since $g_j, h_j \in \mathrm{Pol}_{\leq k-1}(C_n)$ it is easy to see that $\widehat{g}_j, \widehat{h}_j \in \mathrm{Pol}_{\leq k-1}(C_m)$. Equation (4.33) thus shows that $q_{\max} - q$ is $\mathrm{Pol}_{\leq k-1}(C_m)$-sos on C_m. Since q was an arbitrary quadratic form on $m = n/2$ variables, this shows that the $(k-1)$'st level of the SOS hierarchy of $\mathrm{CUT}_{n/2}$ is exact. □

Exponential lower bound If we combine Theorem 18 with Laurent's lower bound on the sum-of-squares hierarchy for the cut polytope (Theorem 17) we obtain the following exponential lower bound for G_{cube}-equivariant SDP lifts of CUT_n.

Theorem 16. *Any G_{cube}-equivariant SDP lift of CUT_n must have size $\geq \binom{n}{\lfloor n/4 \rfloor}$.*

Proof. Let $m = \lfloor n/2 \rfloor$. We apply Theorem 18 with $k = \lceil m/2 \rceil$. Laurent [72] proved that $Q_{k-1}(\mathrm{CUT}_m) \neq \mathrm{CUT}_m$ and thus this means that any G_{cube}-equivariant psd of lift of CUT_n must have size greater than or equal $\binom{n}{k} \geq \binom{n}{\lfloor n/4 \rfloor}$. □

Approximate lifts Like for the parity polytope one can also state a result relating approximate equivariant lifts of the cut polytope and the sum-of-squares hierarchy. To state the result it is convenient to introduce the notion of a (c, s)-*approximation* from the paper [21]. Given two real numbers $c \leq s$ we say that an outer-approximation $\widehat{P}$ of CUT_n achieves a (c, s)-approximation of CUT_n if for any linear form L on $\mathbf{S}^n$ such that $L_{\max} \leq c$ it holds that $\widehat{L_{\max}} \leq s$, where $L_{\max}$ and $\widehat{L_{\max}}$ are respectively the maximum of L on CUT_n and $\widehat{P}$. We can now state the following theorem (we omit the proof since it is very similar to the arguments from the previous proofs):

Theorem 19. *Assume $\widehat{P}$ is an outer-approximation of CUT_n which achieves a (c, s)-approximation and which admits a G_{cube}-equivariant SDP lift of size d. If $d < \binom{n}{k}$ for some $k \leq n/2$ then the $(k-1)$'st sum-of-squares relaxation of $CUT_{\lfloor n/2 \rfloor}$ is a valid (c, s)-approximation of $CUT_{\lfloor n/2 \rfloor}$.*

4.6 Application 3: regular polygons

In this section we study the regular N-gon in the plane and we derive a lower bound on equivariant SDP lifts.

4.6.1 Definitions and symmetry group

The regular N-gon is defined as the convex hull of the N'th roots of unity:

$$\mathcal{X}_N = \{(\cos\theta_k, \sin\theta_k) : k = 0, \ldots, N-1\} \quad \text{where} \quad \theta_k = \frac{2k\pi}{N}.$$

Figure 4-2 shows a picture for $N = 7$. The symmetry group of the regular N-gon is the dihedral group of order $2N$ which consists of N rotations and N reflections. We denote by Rot_N the subgroup of rotations, isomorphic to $\mathbb{Z}_N$.

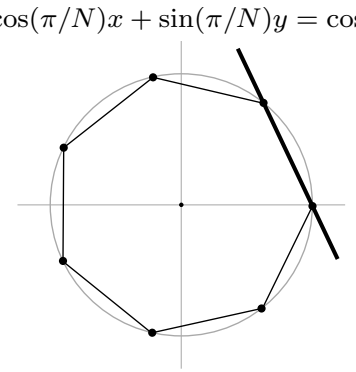

Figure 4-2: The regular 7-gon.

For convenience we will work in this section with Hermitian SDP lifts and complex-valued functions (instead of real-valued). More precisely, a Hermitian SDP lift of a polytope P takes the form $P = \pi(\mathbf{H}^d_+ \cap L)$ where $\mathbf{H}^d_+$ is the cone of $d \times d$ Hermitian positive semidefinite matrices. Also given a nonnegative function f on X a Hermitian sum of squares certificate of f is a certificate

$$f(x) = \sum_{j=1}^{J} |h_j(x)|^2$$

where h_j are complex-valued functions on X.

Our main result in this section is the following.

Theorem 20. *Any* Rot_N*-equivariant Hermitian SDP lift of the regular* N*-gon has size at least* $\ln(N/2)$.

4.6.2 Invariant subspaces of functions on $\mathcal{X}_N$

In order to study Rot_N-equivariant SDP lifts of the regular N-gon, we need to understand the Rot_N-invariant subspaces of $\mathbb{C}^{\mathcal{X}_N}$. As we saw in Example 9, we can think of

$\mathbb{C}^{\mathcal{X}_N}$ as simply the space $\mathbb{C}^{\mathbb{Z}_N}$ where the action of $\mathrm{Rot}_N \cong \mathbb{Z}_N$ shifts the components (since rotating the regular N-gon corresponding to shifting its vertices). We will use this identification in this section.

For $k \in \mathbb{Z}_N$ consider the element $e_k \in \mathbb{C}^{\mathbb{Z}_N}$ defined by:

$$e_k(t) = e^{2ik\pi t/N} \quad \forall t \in \mathbb{Z}_N. \tag{4.34}$$

It is well-known that the functions $(e_k)_{k\in\mathbb{Z}_N}$ form a basis of $\mathbb{C}^{\mathbb{Z}_N}$ which is nothing but the Fourier basis. If $h \in \mathbb{C}^{\mathbb{Z}_N}$, then the decomposition of h in the basis $(e_k)_{k\in\mathbb{Z}_N}$ corresponds to the discrete Fourier transform of h. The $\mathbb{Z}_N$-invariant subspaces of $\mathbb{C}^{\mathbb{Z}_N}$ can be easily expressed in terms of the Fourier basis. This is the object of the next proposition.

Proposition 7. *Any subspace of $\mathbb{C}^{\mathbb{Z}_N}$ that is $\mathbb{Z}_N$-invariant (with respect to the action by shifting) has the form*

$$V = \bigoplus_{k\in K} \mathbb{C}e_k \tag{4.35}$$

where $K \subseteq \mathbb{Z}_N$.

Proof. This is a standard result in Fourier analysis, see e.g., [95, Section 5.1]. We give a short proof for completeness. First note that for any $k \in \mathbb{Z}_N$ the one-dimensional space spanned by e_k is $\mathbb{Z}_N$-invariant since $e_k(t+1) = e^{2ik\pi/N}e_k(t)$ and so $t \mapsto e_k(t+1)$ is in $\mathbb{C}e_k$. It thus follows that any subspace of the form (4.35) is $\mathbb{Z}_N$-invariant. We now show the converse. Assume V is a $\mathbb{Z}_N$-invariant subspace of $\mathbb{C}^{\mathbb{Z}_N}$. We will prove that if $f \in V$ with Fourier decomposition $f = \sum_{k\in\mathbb{Z}_N} \widehat{f}(k)e_k$ then $e_{k_0} \subseteq V$ whenever $\widehat{f}(k_0) \neq 0$. Since V is shift-invariant we have that for any $j \in \mathbb{Z}_N$, $t \mapsto f(t+j)$ is also in V. Thus if $k_0 \in \mathbb{Z}_N$ the function $g_{k_0}(t) = \sum_{j\in\mathbb{Z}_N} f(t+j)e^{-2i\pi jk_0/N}$ is in V. But note that

$$\begin{aligned} g_{k_0}(t) = \sum_{j\in\mathbb{Z}_N} f(t+j)e^{-2i\pi jk_0/N} &= \sum_{j\in\mathbb{Z}_N}\sum_{k\in\mathbb{Z}_N} \widehat{f}(k)e^{2i\pi kt/N}e^{2i\pi kj/N}e^{-2i\pi jk_0/N} \\ &= \sum_{k\in\mathbb{Z}_N} \widehat{f}(k)e^{2i\pi kt/N} \sum_{j\in\mathbb{Z}_N} e^{2i\pi j(k-k_0)/N} \\ &= N\widehat{f}(k_0)e_{k_0}(t). \end{aligned}$$

Thus for $\widehat{f}(k_0) \neq 0$ we get that $e_{k_0} \in V$. □

Facet inequality The "first" facet of the regular N-gon (as shown in Figure 4-2) is defined by the linear inequality

$$\cos(\pi/N) - \cos(\pi/N)x - \sin(\pi/N)y \geq 0.$$

Throughout this section, we denote by ℓ the restriction of this facet on the vertices of the N-gon (which, again, we identify with $\mathbb{Z}_N$):

$$\ell(t) = \cos(\pi/N) - \cos(\pi/N)\cos(2\pi t/N) - \sin(\pi/N)\sin(2\pi t/N), \quad t \in \mathbb{Z}_N. \tag{4.36}$$

The function ℓ has the following expression in the Fourier basis:

$$\ell = \cos(\pi/N)e_0 - \frac{1}{2}e^{-i\pi/N}e_1 - \frac{1}{2}e^{i\pi/N}e_{-1}. \tag{4.37}$$

In order to prove a lower bound on equivariant SDP lifts of the regular N-gon we need to show that ℓ does not admit a sum-of-squares certificate from a low-dimensional invariant subspace. This is made precise in the next theorem, which is just the specialization of the structure theorem (Theorem 12) to the case of the regular N-gon.

Theorem 21 (Structure theorem specialized to regular N-gons)**.** *Assume that the regular N-gon has a Hermitian SDP lift of size d that is equivariant with respect to Rot_N. Then there exists a set $K \subseteq \mathbb{Z}_N$ with $|K| \leq d$ and functions $h_i \in \bigoplus_{k\in K} \mathbb{C}e_k$ such that*

$$\ell = \sum_i |h_i|^2.$$

Proof. We apply Theorem 12 where $G = \mathrm{Rot}_N$ and $x_0 = (1, 0)$ (here the subgroup H that stabilizes x_0 is the trivial one $H = 1_G$). Since G is an abelian group, all the irreducible representations of G have dimension 1, and thus $\alpha_G(d) = 1$ for any d. Thus the theorem says that if $\mathrm{conv}(\mathcal{X}_N)$ has a Rot_N-equivariant SDP lift of size d, then ℓ has a sum-of-squares certificate from an invariant subspace V of $\mathbb{C}^{\mathbb{Z}_N}$ of size d. The theorem then follows from Proposition 7 which states that any such subspace has the form $V = \bigoplus_{k\in K} \mathbb{C}e_k$ where $K \subseteq \mathbb{Z}_N, |K| \leq d$. □

4.6.3 Lower bound on equivariant SDP lifts

This section is dedicated to proving the following theorem:

Theorem 22. *Let $\ell \in \mathbb{C}^{\mathbb{Z}_N}$ be as defined in (4.37) and assume that we can write*

$$\ell = \sum_i |h_i|^2 \quad \text{where} \quad h_i \in \bigoplus_{k\in K} \mathbb{C}e_k \ \ \forall i \tag{4.38}$$

for some set $K \subseteq \mathbb{Z}_N$. Then necessarily $|K| \geq \ln(N/2)$.

This theorem, when combined with Theorem 21, yields the desired lower bound on equivariant SDP lifts of the regular N-gon.

We introduce some notations which will be used throughout the section.

Definition 11. Given $h \in \mathbb{C}^{\mathbb{Z}_N}$ and $K \subseteq \mathbb{Z}_N$, we say that h is *supported on* K and we write $\mathrm{supp}\, h \subseteq K$ if $h \in \bigoplus_{k\in K} \mathbb{C}e_k$.

Definition 12. A set $K \subseteq \mathbb{Z}_N$ is called *sos-valid* if ℓ admits a sum-of-squares certificate $\ell = \sum_i |h_i|^2$ where $\mathrm{supp}\, h_i \subseteq K$ for all i.

Our proof of Theorem 22 proceeds in two steps. In the first step, we give necessary conditions in terms of the "geometry" for a set K to be sos-valid: we show that if

the elements in K can be *clustered* in a certain way then K is not sos-valid. In the second step we propose an algorithm to cluster any given set K, and we prove that our algorithm finds a valid clustering whenever the set K is small enough, i.e., whenever $|K| < \ln(N/2)$.

Necessary conditions for a set to be sos-valid

In this section we give a necessary condition on the "geometry" of a set K to be sos-valid. Before stating the theorem, we make some observations and definitions:

First, observe that if K is a set that is sos-valid, then any *translation* $K' = K + t$ of K is also sos-valid, where $t \in \mathbb{Z}_N$. This is because if $\ell = \sum_i |h_i|^2$ where $\operatorname{supp} h_i \subseteq K$, then we also have $\ell = \sum_i |h'_i|^2$ where $h'_i = e_t h_i$ are supported on K' (since $e_t e_k = e_{t+k}$).

Second, it is useful to think of $\mathbb{Z}_N$ as the nodes of a cycle graph of length N, and of a set of frequencies $K \subseteq \mathbb{Z}_N$ as a subset of the nodes of this graph. For example Figure 4-3 shows a set K with $|K| = 7$ for the $N = 12$-gon (the elements of K are the black dots). Note that since the property of being sos-valid is invariant under translation, the cycle graph need not be labeled. The only information that matters are the relative distances of the elements of K with respect to each other.

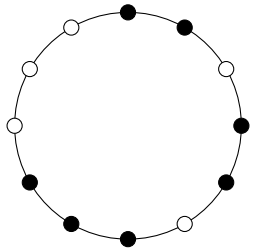

Figure 4-3: A set of frequencies K for the regular 12-gon.

We endow $\mathbb{Z}_N$ with the natural distance d on the cycle graph. The distance between two frequencies $k, k' \in \mathbb{Z}_N$ is denoted by $d(k, k')$; also if C, C' are two subsets of $\mathbb{Z}_N$ we let

$$d(C, C') = \min_{k \in C, k' \in C'} d(k, k').$$

If $x \in \mathbb{Z}_N$ and r is a positive integer, we can define the ball $B(x, r)$ centered at x and with radius r to be the set $B(x, r) := \{y \in \mathbb{Z}_N : d(x, y) \leq r\}$. We also let $[x, x + r]$ be the interval $\{x, x + 1, \ldots, x + r\} \subseteq \mathbb{Z}_N$. Note that the ball centered at x of radius r is simply the interval $[x - r, x + r]$.

To state the main theorem, it will be more convenient to work with diameters instead of radii of balls (mainly to avoid the issue of dividing by two). We introduce the notion of *in-diameter* of a set K which is essentially twice the radius of the smallest ball containing K. More formally we have:

Definition 13. Let $K \subseteq \mathbb{Z}_N$. We define the *in-diameter* of K, denoted $\mathbf{diam}_{\text{in}}(K)$ to be the smallest positive integer r such that K is included in an interval of the form $[x, x + r]$ where $x \in \mathbb{Z}_N$.

Remark 9. Note that the in-diameter of a set K is in general different from the usual notion of *diameter* (largest distance between two elements in K). Note for example that $\mathbf{diam}_{\text{in}}(\mathbb{Z}_N) = N$ whereas the diameter of $\mathbb{Z}_N$ is equal to $\lfloor N/2 \rfloor$.

We are now ready to state the main result of this section:

Theorem 23. *Let N be an integer and let $K \subseteq \mathbb{Z}_N$ be a set of frequencies. Assume that K can be decomposed into disjoint clusters $(C_\alpha)_{\alpha \in A}$:*

$$K = \bigcup_{\alpha \in A} C_\alpha,$$

such that the following holds for some $1 \leq \gamma < N/2$:

(i) For any $\alpha \in A$, C_α has in-diameter $\leq \gamma$.

(ii) For any $\alpha \neq \alpha'$, $d(C_\alpha, C_{\alpha'}) > \gamma$.

Then the set K is not *sos-valid (i.e., it is* not *possible to write the linear function ℓ as a sum of squares of functions supported on K).*

Proof. To prove this theorem, we will construct a linear functional $\mathcal{L}$ on $\mathbb{C}^{\mathbb{Z}_N}$ such that:

(a) $\mathcal{L}(\ell) < 0$, and;

(b) for any h supported on K we have $\mathcal{L}(|h|^2) \geq 0$.

Clearly this will show that we cannot have $\ell = \sum_i |h_i|^2$ where $\operatorname{supp} h_i \subseteq K$.

We first introduce a piece of notation that will be needed for the definition of $\mathcal{L}$: Given $k \in \mathbb{Z}_N$, we let $k \bmod N$ be the unique element in

$$\Big\{ -\lceil N/2 \rceil + 1, \ldots, \lfloor N/2 \rfloor \Big\}$$

that is equal to k modulo N. The main property that will be used about this operation is the following, which can be verified easily: If $k, k' \in [0, \gamma]$ where $\gamma < N/2$ then:

$$(k' - k) \bmod N = (k' \bmod N) - (k \bmod N). \tag{4.39}$$

Let $p = e^{i\pi/N}$ and note that p does not belong to our regular N-gon. We now define the linear functional $\mathcal{L} : \mathbb{C}^{\mathbb{Z}_N} \to \mathbb{C}$ as follows, for all $k \in \mathbb{Z}_N$:

$$\mathcal{L}(e_k) = \begin{cases} p^{k \bmod N} & \text{if } d(0,k) \leq \gamma \\ 0 & \text{else.} \end{cases} \tag{4.40}$$

Note that the map $\mathcal{L}$ can be interpreted as the composition of two maps:

$$\mathcal{L} = \mathsf{Eval}_p \circ \mathcal{E}$$

where $\mathcal{E}$ is a map that extrapolates a function $h \in \mathbb{C}^{\mathbb{Z}_N}$ defined on the vertices of the N-gon to a function on the unit circle, and Eval_p is a map that evaluates a function on the unit circle to the point p. The extrapolation map $\mathcal{E}$ is defined on the Fourier basis as follows: $\mathcal{E}(e_k)(z) = z^{k \bmod N}$ if $d(0,k) \leq \gamma$ and 0 otherwise, for z in the unit circle.

We now prove that $\mathcal{L}$ satisfies properties (a) and (b) above.

(a) It is easy to see that $\mathcal{L}(\ell) < 0$. Indeed since $\gamma \geq 1$ we have $\mathcal{L}(e_1) = e^{i\pi/N}$ and $\mathcal{L}(e_{-1}) = e^{-i\pi/N}$ which implies that:

$$\mathcal{L}(\ell) = \mathcal{L}\Big(\cos(\pi/N)e_0 - (e^{-i\pi/N}e_1 + e^{i\pi/N}e_{-1})/2\Big) = \cos(\pi/N) - 1 < 0.$$

(b) We now show that if h is a function supported on K, then $\mathcal{L}(|h|^2) \geq 0$. Since $K = \cup_{\alpha \in A} C_\alpha$, we can write

$$h = \sum_{k \in K} h_k e_k = \sum_{\alpha \in A} \sum_{k \in C_\alpha} h_k e_k.$$

Thus

$$|h|^2 = h^* h = \underbrace{\sum_{\alpha \in A} \left| \sum_{k \in C_\alpha} h_k e_k \right|^2}_{P} + \underbrace{\sum_{\alpha \neq \alpha'} \sum_{k \in C_\alpha, k' \in C_{\alpha'}} h_k^* h_{k'} e_k^* e_{k'}}_{Q}. \tag{4.41}$$

Let P and Q be the first and second terms in the equation above. We will show that $\mathcal{L}(Q) = 0$ and that $\mathcal{L}(P) \geq 0$. Observe that if $k \in C_\alpha$ and $k' \in C_{\alpha'}$ where $\alpha \neq \alpha'$ then we have:

$$\mathcal{L}(e_k^* e_{k'}) = \mathcal{L}(e_{k'-k}) = 0$$

where the last equality follows since $d(k'-k, 0) = d(k', k) > \gamma$ (cf. assumption (ii) on the clustering). Thus this shows that $\mathcal{L}(Q) = 0$.

We will now show that $\mathcal{L}(P) \geq 0$, by showing that for any $\alpha \in A$ we have

$$\mathcal{L}\left(\left| \sum_{k \in C_\alpha} h_k e_k \right|^2 \right) \geq 0.$$

Let $\alpha \in A$. By assumption (i) on the clustering, we know that the in-diameter of C_α is $\leq \gamma$, i.e., that C_α is included in an interval $[x, x+\gamma]$. Note that since

$$\left| \sum_{k \in C_\alpha} h_k e_k \right|^2 = \left| e_{-x} \sum_{k \in C_\alpha} h_k e_k \right|^2 = \left| \sum_{k \in C_\alpha} h_k e_{k-x} \right|^2$$

we can assume without loss of generality that $x = 0$. Now since $C_\alpha \subseteq [0, \gamma]$, we

have from (4.39) that for any $k, k' \in C_\alpha$:

$$(k' - k) \bmod N = (k' \bmod N) - (k \bmod N) \tag{4.42}$$

Using this we have:

$$\begin{aligned}\mathcal{L}\left(\left|\sum_{k\in C_\alpha} h_k e_k\right|^2\right) = \sum_{k,k'\in C_\alpha} h_k^* h_{k'} \mathcal{L}(e_{k'-k}) &\overset{(a)}{=} \sum_{k,k'\in C_\alpha} h_k^* h_{k'} p^{(k'-k) \bmod N} \\ &\overset{(b)}{=} \sum_{k,k'\in C_\alpha} h_k^* h_{k'} p^{k' \bmod N} p^{-(k \bmod N)} \\ &= \left|\sum_{k\in C_\alpha} h_k p^{k \bmod N}\right|^2 \geq 0\end{aligned}$$

where in (a) we used the fact that $d(0, k'-k) = d(k', k) \leq \gamma$ and in (b) we used identity (4.42). Thus this shows that $\mathcal{L}(|h|^2) \geq 0$ for all h supported on C_α, which implies that $\mathcal{L}(P) \geq 0$ (since $P = \sum_{\alpha\in A} \left|\sum_{k\in C_\alpha} h_k e_k\right|^2$) which is what we wanted.

□

Remark 10. To illustrate the previous theorem consider the following two simple applications:

- It is not hard to prove that the theta-rank of the regular N-gon is at least $N/4$ (see e.g., [36]). This lower bound can be obtained as a direct corollary of Theorem 23. Indeed if K is contained in the open interval $(-\lceil N/4\rceil, \lceil N/4\rceil)$, then the in-diameter of K is $< N/2$ which means that if we consider K as a single cluster, it satisfies conditions (i) and (ii) of the theorem with $\gamma = \mathbf{diam}_{\text{in}}(K)$. Thus such a K is not sos-valid.

- We can also give another simple application of the previous theorem: Assume K is a set of frequencies that has no two consecutive frequencies, i.e., for any $k, k' \in K$ where $k \neq k'$ we have $d(k, k') \geq 2$. It is not hard to see that such a set K cannot be sos-valid: indeed if h is a function supported on K, then the expansion of $|h|^2$ does not have any term involving the frequencies e_1 or e_{-1}. Thus it is not possible to write ℓ as a sum-of-squares of elements supported on such K. This simple fact can be obtained as a consequence of Theorem 23 if we consider each frequency of K as its own cluster (i.e., we write $K = \cup_{k\in K}\{k\}$) and conditions (i) and (ii) of the theorem are satisfied with $\gamma = 1$.

An algorithm to find valid clusterings and a logarithmic lower bound

We now study sets K which admit a clustering that satisfies points (i) and (ii) of Theorem 23. The main purpose of this section is to show that any set K with $|K| < \ln(N/2)$ admits such a clustering, which implies that it cannot be sos-valid.

This would thus show that any Rot_N-equivariant Hermitian SDP lift of the regular N-gon has to have size at least $\ln(N/2)$.

For convenience we call a *valid clustering* of a set K, any clustering that satisfies points (i) and (ii) of Theorem 23. We state this in the following definition for future reference:

Definition 14. Let $K \subseteq \mathbb{Z}_N$. We say that K has a *valid clustering* if K can be decomposed into disjoint clusters $(C_\alpha)_{\alpha \in A}$:

$$K = \bigcup_{\alpha \in A} C_\alpha,$$

such that the following holds for some $1 \leq \gamma < N/2$:

(i) For any $\alpha \in A$, C_α has in-diameter $\leq \gamma$.

(ii) For any $\alpha \neq \alpha'$, $d(C_\alpha, C_{\alpha'}) > \gamma$.

We propose a simple greedy algorithm to search for a valid clustering for any set $K \subseteq \mathbb{Z}_N$: We start with each point of K in its own cluster and at each iteration we merge the two closest clusters. We keep doing this until we get a clustering that satisfies the required condition, or until all the points are in the same cluster. We show in this section that if the number of points of K is small enough, if $|K| < \ln(N/2)$, then this algorithm terminates by producing a valid clustering of K. For reference we describe the algorithm more formally in Algorithm 1.

Algorithm 1 Algorithm to produce a clustering of a set K

Input: A set $K \subseteq \mathbb{Z}_N$

Output: A *valid clustering* of K (in the sense of Definition 14) or "0" if no valid clustering found.

• Consider initial clustering where each element of K is in its own cluster. If this clustering is already valid (which is equivalent to say that for any distinct elements $k, k' \in K$ we have $d(k, k') \geq 2$) then output this clustering as a valid clustering with parameter $\gamma = 1$.

• Precompute the pairwise distances between points in K and sort these distances in increasing order $d_1 \leq d_2 \leq d_3 \leq \ldots$ (cf. Figure 4-4; note that two distances d_i and d_j could be equal).

for $i = 1, 2, \ldots, |K| - 1$ **do**

Let $x, y \in K$ be the i'th closest points in K so that $d(x, y) = d_i$. If x and y are in different clusters, then merge these two clusters.

If the current clustering satisfies points (i) and (ii) of Definition 14 (with γ equal to the largest in-diameter in all the clusters) stop and output the current clustering.

end for

If no valid clustering was found, output "0"

In the next theorem, we show that any set $K \subseteq \mathbb{Z}_N$ with $|K| < \ln(N/2)$ has a valid clustering.

Theorem 24. *If a set $K \subseteq \mathbb{Z}_N$ satisfies $|K| < \ln(N/2)$, then a valid clustering of K exists and Algorithm 1 will produce one.*

Proof. Observe that at the end of iteration i of the algorithm, the distance between any pair of clusters is greater than or equal d_{i+1}: Assume for contradiction that there are two clusters C, C' at iteration i where $d(C, C') < d_{i+1}$. This means that there exist $x \in C$, $y \in C'$ such that $d(x, y) < d_{i+1}$. But this is impossible because the algorithm processes distances in increasing order, and so x and y must have merged in the same cluster at some iteration $\leq i$.

Now, to prove that the algorithm terminates and produces a valid clustering, we need to show that at some iteration i, each cluster has in-diameter smaller than $\min(d_{i+1}, N/2)$. Note that one can get a simple upper bound on the in-diameter of the clusters at iteration i: indeed, it is not hard to show that at iteration i any cluster has in-diameter at most S_i, where S_i is defined as:

$$S_i := d_1 + d_2 + \cdots + d_i = \sum_{j=1}^{i} d_j.$$

Figure 4-4 shows a simple illustration of this bound.

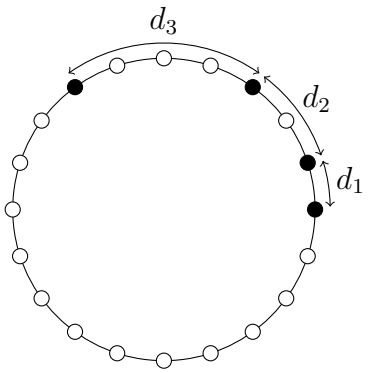

Figure 4-4: A set of frequencies K. At iteration 0 of the algorithm each frequency is in its own cluster. At iteration 1 of the algorithm, the two nodes at distance d_1 from each other are merged in a single cluster. At iteration 2, the two nodes at distance d_2 are merged and we get one cluster having 3 nodes with in-diameter $d_1 + d_2$. In general, at iteration i the clusters cannot have in-diameter larger than $d_1 + \cdots + d_i$.

Let a be the largest index i where $d_i = 1$, and let b the largest index i where $S_i < N/2$.[2] If $i \in [a, b]$, then at the end of the i'th iteration, the distance between any two clusters is greater than 1 (since $d_{i+1} > 1$) and the in-diameter of any cluster is smaller than $N/2$. To prove that the algorithm terminates and produces a valid clustering, it suffices to show that there exists $i \in [a, b]$ such that $d_{i+1} > S_i$.

[2]Note that we can assume $\mathbf{diam}_{\text{in}}(K) \geq N/2$ which implies that $S_{|K|-1} \geq N/2$. Indeed, if the in-diameter of K is smaller than $N/2$, then we have a valid clustering of K by considering K as a single cluster.

Assume for contradiction that this is not the case. Then this means that we have:

$$\begin{aligned} d_{a+1} &\leq d_1 + \cdots + d_a \\ d_{a+2} &\leq d_1 + \cdots + d_{a+1} \\ &\vdots \\ d_{b+1} &\leq d_1 + \cdots + d_b \end{aligned}$$

We will now show that this implies that $|K| \geq \ln(N/2)$ which contradicts the assumption of the theorem. Define the function $f(x) = 1/x$ and note that, on the one hand we have:

$$\sum_{i=a}^{b} d_{i+1} f(S_i) = \sum_{i=a}^{b} d_{i+1} \frac{1}{d_1 + \cdots + d_i} \leq \sum_{i=a}^{b} 1 = b - a + 1.$$

On the other hand, since f is a decreasing function we have (cf. Figure 4-5):

$$\sum_{i=a}^{b} d_{i+1} f(S_i) \geq \int_{S_a}^{S_{b+1}} f(x)dx = [\ln(x)]_{S_a}^{S_{b+1}} = \ln(S_{b+1}) - \ln(S_a).$$

Thus we get that:

$$b - a + 1 \geq \ln(S_{b+1}) - \ln(S_a).$$

Now note that $S_a = a$ since $d_i = 1$ for all $1 \leq i \leq a$. Thus we have:

$$b \geq \ln(S_{b+1}) - \ln(S_a) + a - 1 \geq \ln(S_{b+1})$$

since $a - \ln(S_a) \geq 1$ (we assume here that $a \geq 1$ because otherwise the distance between any two elements in K is at least 2 in which case K is clearly not sos-valid). Now since $|K| \geq b$ and $S_{b+1} \geq N/2$ we get

$$|K| \geq \ln(N/2)$$

as desired.

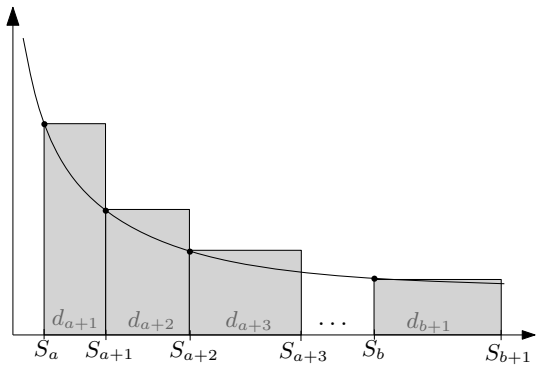

Figure 4-5:

□

4.7 Summary of chapter

- Given a polytope $P = \text{conv}(X)$ we saw in Theorem 5 (Chapter 2) that a small SDP lift of P can be obtained by finding a low-dimensional subspace V of $\mathbb{R}^X$ such that any facet inequality of P can be written as a sum of squares of functions from V.
- If P has symmetries and is invariant under the action of a group G, then a sufficient condition to get an SDP lift that "respects" this symmetry is to require the subspace V to be G-invariant (see Theorem 9). Under some additional mild conditions this is also necessary (see Theorem 10). The necessity part is what we called the "structure theorem" for equivariant SDP lifts.
- To prove lower bounds on equivariant SDP lifts of $P = \text{conv}(X)$ one has to understand the structure of G-invariant subspaces of $\mathbb{R}^X$. Theorems 13 and 16 establish exponential lower bounds on equivariant SDP lift of the parity polytope and the cut polytope. Theorem 20 shows that any equivariant SDP lift of the regular N-gon must have size $\Omega(\log N)$.

4.8 Proofs

4.8.1 Proof of Theorem 9: equivariance of sum of squares lifts when subspace V is G-invariant

Let X be a finite set in $\mathbb{R}^n$ and assume X is invariant under the action of a group $G \subseteq GL_n(\mathbb{R})$. We also assume that $\text{conv}(X)$ is full-dimensional. In this appendix, we show that the SDP lift (2.20) is G-equivariant when the subspace V is chosen to be G-invariant.

Assume V is a subspace of $\mathbb{R}^X$ such that any valid linear inequality on $\text{conv}(X)$ has a sum-of-squares certificate from V. Then we saw in Section 2.2.5 (cf. Equation (2.20)) that we have the following description of $\text{conv}(X)$:

$$\text{conv}(X) = \Big\{ (E(e_1), \ldots, E(e_n)) \ : E \in (\mathbb{R}^X)^* \text{ where } E(1) = 1,\ E(f^2) \geq 0 \ \forall f \in V \Big\}. \tag{4.43}$$

where for each i, $e_i \in \mathbb{R}^X$ is the function defined by $e_i(x) = x_i$. Equation (4.43) expresses $\text{conv}(X)$ as a SDP lift of size $d = \dim V$. Indeed the last constraint on E

in (4.43) is equivalent to saying that the bilinear form H_E on V defined by

$$H_E : V \times V \to \mathbb{R}, \quad H_E(f_1, f_2) = E(f_1 f_2)$$

is positive semidefinite. Thus if we identify $\mathbf{S}^d$ with bilinear forms on V (by fixing a basis) then Equation (4.43) can be rewritten as:

$$\text{conv}(X) = \pi(\mathbf{S}^d_+ \cap L) \tag{4.44}$$

where

- $L \subset \mathbf{S}^d$ is the affine subspace

$$L := \{H_E \ : \ E \in (\mathbb{R}^X)^*, E(1) = 1\},$$

  i.e., L the image of the affine space $\{E \in (\mathbb{R}^X)^*, E(1) = 1\}$ under the linear map $E \mapsto H_E$;

- and π is the linear map, which given a bilinear form of the form H_E for some $E \in (\mathbb{R}^X)^*$, returns the n-tuple $(E(e_1), \ldots, E(e_n)) \in \mathbb{R}^n$ (this linear map π is well-defined when P is full-dimensional, since one can show in this case that the functions e_i are all in $\text{span}\{f^2 : f \in V\}$).

We now proceed to show that the SDP lift (4.44) satisfies the definition of G-equivariance, where G is the automorphism group of X.

Since G acts on $\mathbb{R}^X$, it also acts on the dual space $(\mathbb{R}^X)^*$ as follows: If $E \in (\mathbb{R}^X)^*$ then we let

$$(g \cdot E)(f) := E(g^{-1} \cdot f) \quad \forall f \in \mathbb{R}^X.$$

Equivariance of the lift (4.44) now follows from the following main lemma:

Lemma 5. *Given $E \in (\mathbb{R}^X)^*$ we have for any $g \in G$ and any $f, h \in V$:*

$$H_{g \cdot E}(f_1, f_2) = H_E(g^{-1} \cdot f_1, g^{-1} \cdot f_2). \tag{4.45}$$

Remark 11. One can interpret the identity (4.45) in matrix terms as follows: Given $g \in G$, let $\theta(g)$ be the $d \times d$ matrix which corresponds to the linear map $f \in V \mapsto g \cdot f$. Define $\rho(g) = \theta(g^{-1})^T$. Then identity (4.45) is the same as:

$$H_{g \cdot E} = \rho(g) H_E \rho(g)^T \tag{4.46}$$

where $H_{g \cdot E}$ and H_E are interpreted as symmetric matrices of size d.

Proof. The following sequence of equalities proves the claim:

$$\begin{aligned} H_{g \cdot E}(f_1, f_2) = (g \cdot E)(f_1 f_2) &= E(g^{-1} \cdot (f_1 f_2)) \\ &\overset{(*)}{=} E((g^{-1} \cdot f_1)(g^{-1} \cdot f_2)) = H_E(g^{-1} \cdot f_1, g^{-1} \cdot f_2). \end{aligned}$$

In Equality (*) we used the fact that the action of G on $\mathbb{R}^X$ satisfies:

$$g \cdot (f_1 f_2) = (g \cdot f_1)(g \cdot f_2)$$

which can be easily seen since for any $x \in X$ we have:

$$g \cdot (f_1 f_2)(x) = (f_1 f_2)(g^{-1} \cdot x) = f_1(g^{-1} \cdot x) f_2(g^{-1} \cdot x) = (g \cdot f_1)(x)(g \cdot f_2)(x).$$

□

Using this lemma it is easy to check that the lift (4.44) satisfies the definition of G-equivariance.

4.8.2 Proof of Theorem 11: factorization theorem for equivariant SDP lifts

Proof of Theorem 11. We first treat the case where the stabilizer of x_0 is $\{1_G\}$. In this case the proof is almost trivial: Let $A(x_0)$ be any point in $\mathbf{S}^d_+ \cap L$ such that $\pi(A(x_0)) = x_0$. Define, for any $g \in G$, $A(g \cdot x_0) := \rho(g) A(x_0) \rho(g)^T$. Note that $A(g \cdot x_0) \in \mathbf{S}^d_+ \cap L$ by the definition of an equivariant SDP lift (Definition 9). Also note that

$$\pi(A(g \cdot x_0)) = \pi(\rho(g) A(x_0) \rho(g)^T) \overset{(a)}{=} g \cdot \pi(A(x_0)) \overset{(b)}{=} g \cdot x_0 \tag{4.47}$$

where in (a) we used the definition of equivariant SDP lift, and in (b) we used the fact that $\pi(A(x_0)) = x_0$. Since $X = G \cdot x_0$, Equation (4.47) shows that $\pi(A(x)) = x$ for all $x \in X$. Thus we can use Theorem 3 (cf. Remark 4) with our choice of map A to show that Property (i) of the statement holds. Note that Property (ii) holds by construction of the map A.

We now treat the general case. Let H be the stabilizer of x_0. Let A_0 be any point in $\mathbf{S}^d_+ \cap L$ such that $\pi(A_0) = x_0$ and define:

$$A(x_0) = \frac{1}{|H|} \sum_{h \in H} \rho(h) A_0 \rho(h)^T.$$

Now we extend the map A to the whole X by letting $A(x) = \rho(g) A(x_0) \rho(g)^T$ where g is any element of G such that $g \cdot x_0$. Note that this is well-defined since if $g \cdot x_0 = g' \cdot x_0$, we have, with $h = g^{-1} g' \in H$:

$$\begin{aligned} \rho(g') A(x_0) \rho(g')^T &= \rho(gh) A(x_0) \rho(gh)^T \\ &= \rho(g) \rho(h) A(x_0) \rho(h)^T \rho(g)^T = \rho(g) A(x_0) \rho(g)^T. \end{aligned}$$

It is easy to verify, like in the previous case, that the map we just defined satisfies $\pi(A(x)) = x$ for all $x \in X$. Thus by applying Theorem 3 with our choice of map A we get that Property (i) of the statement holds. Also Property (ii) holds by construction of the map A.

Note that we can assume $\rho(g)$ to be orthogonal for any $g \in G$, by a simple change of basis: By Proposition 4 let Q be an invertible matrix such that $Q\rho(g)Q^{-1}$ is orthogonal. By letting $\widehat{\rho}(g) = Q\rho(g)Q^{-1}$, $\widehat{A}(x) = QA(x)Q^T$, $\widehat{B} = Q^{-T}BQ^{-1}$ we have $\langle \widehat{A}(x), \widehat{B}\rangle = \langle A(x), B\rangle$, $\widehat{A}(g \cdot x) = \widehat{\rho}(g)\widehat{A}(x)\widehat{\rho}(g)^T$ and $\widehat{\rho}(g)\widehat{\rho}(g)^T = I_d$ which is what we need. □

4.8.3 Proof of Lemma 3: irreducible subspaces of $\mathbb{R}^{\mathbf{EVEN}_n}$

Proof of Lemma 3. It is easy to see that $\mathrm{Pol}_k(\mathrm{EVEN}_n)$ is invariant for each $k = 0, 1, \ldots, \lfloor n/2 \rfloor$. It remains to show that each of these is irreducible. For any $k \neq \ell \in [n]$, let $\epsilon_{k,\ell} \in G_{\text{parity}}$ be defined by $\epsilon_{k,\ell} = \mathrm{diag}(1, \ldots, -1, \ldots, -1, \ldots, 1)$ where all the entries are equal to 1 except the entries in position k and ℓ which are equal to -1. Given an element $p \in \mathbb{R}^{\mathrm{EVEN}_n}$ we denote by $(\mathrm{id} + \epsilon_{k\ell}) \cdot p$ the polynomial $p + \epsilon_{k\ell} \cdot p$. Observe that whenever $I \subset [n]$, then

$$(\mathrm{id} + \epsilon_{k\ell}) \cdot x^I = x^I + \epsilon_{k\ell} \cdot x^I = \begin{cases} 2x^I & \text{if } (k \in I \text{ and } \ell \in I) \text{ or } (k \notin I \text{ and } \ell \notin I) \\ 0 & \text{if either exactly one of } k \in I \text{ and } \ell \in I \text{ occur.} \end{cases}$$

Now fix some arbitrary $k < n/2$ (we deal with the case $n = 2k$ separately) and let V be a (non-zero) invariant subspace of $\mathrm{Pol}_k(\mathrm{EVEN}_n)$. Let p be a non-zero element of V. By the invariance of V under the permutation action, we can assume that the coefficient of the monomial $x_1x_2\cdots x_k$ in p is non-zero, and so p is of the form: $p(x) = cx_1\cdots x_k + \sum_{|I|=k, I\neq\{1,2,\ldots,k\}} c_I x^I$ where $c \neq 0$. We will show that necessarily $V = \mathrm{Pol}_k(\mathrm{EVEN}_n)$. We first show that

$$\left[\prod_{i=k+2}^{n} (\mathrm{id} + \epsilon_{k+1,i})\right] \cdot p(x) = 2^{n-k-1} x_1x_2\cdots x_k. \tag{4.48}$$

Once this is established we will know, by the $\mathfrak{S}_n$-invariance of V, that V is equal to $\mathrm{Pol}_k(\mathrm{EVEN}_n)$.

To establish (4.48), first note that if $i \in \{k+2, k+3, \ldots, n\}$ then

$$(\mathrm{id} + \epsilon_{k+1,i})(x_1x_2\cdots x_k) = 2x_1x_2\cdots x_k$$

because neither of x_i and x_{k+1} appear in $x_1x_2\cdots x_k$. It remains to check that every other monomial of degree k is in the kernel of $\left[\prod_{i=k+2}^{n}(\mathrm{id} + \epsilon_{k+1,i})\right]$. Consider any other monomial x^I, i.e. $I \subset \{1, 2, \ldots, n\}$ with $|I| = k$ and for which there is some $\ell \in I$ with $\ell \geq k+1$. Consider two cases, first the case where $k+1 \notin I$. Then there is some $\ell \geq k+2$ such that $\ell \in I$. But then $(\mathrm{id} + \epsilon_{k+1,\ell}) \cdot x^I = 0$ and so since the $\epsilon_{i,j}$ commute, $\left[\prod_{i=k+2}^{n}(\mathrm{id} + \epsilon_{k+1,i})\right] \cdot x^I = 0$. Now suppose $k+1 \in I$. Then there is some $\ell \geq k+2$ such that $\ell \notin I$. This is because if there were no such ℓ then we must have $I \supseteq \{k+1, k+2, \ldots, n\}$ which cannot have cardinality k since we assumed $k < n/2$. It then follows that $(\mathrm{id} + \epsilon_{k+1,\ell}) \cdot x^I = 0$ and so that $\left[\prod_{i=k+2}^{n}(\mathrm{id} + \epsilon_{k+1,i})\right] \cdot x^I = 0$.

Finally consider the case when $n = 2k$. In this situation since V is invariant under

the permutation action we can assume

$$p(x) = c(x_1 \cdots x_{n/2} + x_{n/2+1} \cdots x_n) + \sum_{\substack{|I|=n/2, I \neq \{1,2,\ldots,n/2\} \\ I \neq \{n/2, n/2+1, \ldots, n\}}} c_I(x^I + x^{I^c}).$$

Applying the same argument as above, we see that

$$\left[\prod_{i=k+2}^{n} (\mathrm{id} + \epsilon_{k+1,i})\right] \cdot p(x) = 2^{n-n/2-1} \left(x_1 x_2 \cdots x_{n/2} + x_{n/2+1} x_{n/2+2} \cdots x_n\right).$$

Since the action of $\mathfrak{S}_n$ on

$$x_1 x_2 \cdots x_{n/2} + x_{n/2+1} x_{n/2+2} \cdots x_n$$

generates a basis for $\mathrm{Pol}_{n/2}(\mathrm{EVEN}_n)$ we can conclude that $V = \mathrm{Pol}_{n/2}(\mathrm{EVEN}_n)$.

The second part of the theorem is a direct consequence of Proposition 3 on low-dimensional invariant subspaces. When n is odd note that $\dim \mathrm{Pol}_0(\mathrm{EVEN}_n) < \dim \mathrm{Pol}_1(\mathrm{EVEN}_n) < \cdots < \dim \mathrm{Pol}_{\lfloor n/2 \rfloor}(\mathrm{EVEN}_n)$ with $\dim \mathrm{Pol}_k(\mathrm{EVEN}_n) = \binom{n}{k}$. Thus any invariant subspace V of $\mathbb{R}^{\mathrm{EVEN}_n}$ with $\dim V < \dim \mathrm{Pol}_k(\mathrm{EVEN}_n) = \binom{n}{k} = D_{n,k}$ must be contained in $V_{k-1} = \mathrm{Pol}_0(\mathrm{EVEN}_n) \oplus \cdots \oplus \mathrm{Pol}_{k-1}(\mathrm{EVEN}_n)$ and thus consists of polynomials of degree at most $k-1$. In the case where n is even we have $\dim \mathrm{Pol}_{n/2}(\mathrm{EVEN}_n) = \frac{1}{2}\binom{n}{n/2}$. Thus any invariant subspace V with $\dim V < \min\left(\binom{n}{k}, \frac{1}{2}\binom{n}{n/2}\right) = D_{n,k}$ must be contained in $\mathrm{Pol}_{\leq k-1}(\mathrm{EVEN}_n) = \mathrm{Pol}_0(\mathrm{EVEN}_n) \oplus \cdots \oplus \mathrm{Pol}_{k-1}(\mathrm{EVEN}_n)$. □

4.8.4 Proof of Proposition 6: lower bound on theta rank of parity polytope

Before proving the lower bound on the theta rank, we need the following simple lemma which expresses the pointwise product of the functions e_I and e_J in terms of our basis for $\mathbb{R}^{\mathrm{EVEN}_n}$. (Here, and subsequently, if $I, J \subseteq [n]$ then $I \triangle J$ is the symmetric difference of I and J.)

Lemma 6. *If $I, J \subseteq [n]$ then*

$$e_I e_J = \begin{cases} e_{I \triangle J} & \text{if } |I \triangle J| < n/2 \\ e_{(I \triangle J)^c} & \text{if } |I \triangle J| > n/2 \\ (e_{I \triangle J} + e_{(I \triangle J)^c})/2 & \text{if } |I \triangle J| = n/2. \end{cases}$$

Proof. For any $x \in \mathrm{EVEN}_n$ we have that $x_i^2 = 1$ for all $i \in [n]$ and $\prod_{i \in [n]} x_i = 1$. Hence for any $I, J \subseteq [n]$, $x^I x^J = x^{I \triangle J} = x^{(I \triangle J)^c} = (x^{I \triangle J} + x^{(I \triangle J)^c})/2$. The result then follows by recognizing that at least one of these can be written in terms of the given basis from Proposition 5. □

Proof of Proposition 6. We first sketch the outline of the proof before filling in the details. We choose a function $\ell_{\max} - \ell \in \mathrm{Pol}_{\leq 1}(\mathrm{EVEN}_n)$ that is nonnegative on EVEN_n but when viewed as a polynomial on $\mathbb{R}^n$ takes a negative value on some point $p \in \{-1,1\}^n \setminus \mathrm{EVEN}_n$. (One can think of this as a facet inequality for PAR_n that is not valid for the hypercube.) We use this point p to construct a linear functional $\mathcal{L}_p \in (\mathbb{R}^{\mathrm{EVEN}_n})^*$ (defined to mimic evaulation of the function at p) such that $\mathcal{L}_p(\ell_{\max} - \ell) < 0$ and yet whenever $k < n/4$ and $f \in \mathrm{Pol}_{\leq k}(\mathrm{EVEN}_n)$ we have that $\mathcal{L}_p(f^2) \geq 0$. This would imply that if $k < n/4$ then $\mathcal{L}_p$ separates the linear function $\ell_{\max} - \ell$ (that is nonnegative on EVEN_n) from the cone of functions that are $\mathrm{Pol}_{\leq k}(\mathrm{EVEN}_n)$-sos on EVEN_n, which would complete the proof.

We now fill in the details. Let $\ell_{\max} - \ell(x) = (n-2)e_\emptyset(x) + e_{\{n\}}(x) - \sum_{i=1}^{n-1} e_{\{i\}}(x)$ and observe that this is a nonnegative function on EVEN_n (since it defines a facet of PAR_n). Let $p = (1,1,\ldots,1,-1) \in \mathbb{R}^n$, and note that p has an *odd* number of -1s. Define a linear functional $\mathcal{L}_p \in (\mathbb{R}^{\mathrm{EVEN}_n})^*$ by defining it on our basis for $\mathbb{R}^{\mathrm{EVEN}_n}$ by

$$\mathcal{L}_p(e_I) = \prod_{i \in I} p_i \text{ if } |I| < n/2 \text{ and}$$
$$\mathcal{L}_p((e_I + e_{I^c})/2) = \left(\prod_{i \in I} p_i + \prod_{i \in I^c} p_i \right)/2 = 0 \text{ if } |I| = n/2.$$

Observe that $\mathcal{L}_p(\ell_{\max} - \ell) = (n-2) + \mathcal{L}_p(e_{\{n\}}) - \sum_{i=1}^{n-1} \mathcal{L}_p(e_{\{i\}}) = (n-2) - 1 - (n-1) = -2$.

We now show that if $k < n/4$ and $f \in \mathrm{Pol}_{\leq k}(\mathrm{EVEN}_n)$ then $\mathcal{L}_p(f^2) \geq 0$. Observe that if $|I \triangle J| < n/2$ and $|I|, |J| < n/2$ then

$$\mathcal{L}_p(e_I e_J) = \mathcal{L}_p(e_{I \triangle J}) = \prod_{i \in I \triangle J} p_i = \left(\prod_{i \in I} p_i\right)\left(\prod_{j \in J} p_j\right) = \mathcal{L}_p(e_I)\mathcal{L}_p(e_J). \tag{4.49}$$

If $f \in \mathrm{Pol}_{\leq k}(\mathrm{EVEN}_n)$ where $k < n/4$ then there are constants $c_I \in \mathbb{R}$ such that $f = \sum_{m=0}^{k} \sum_{|I|=k} c_I e_I$. Hence by (4.49), and the observation that if $|I|, |J| < n/4$ then $|I \triangle J| < n/2$, we can see that any $f \in \mathrm{Pol}_{\leq k}(\mathrm{EVEN}_n)$ satisfies

$$\begin{aligned} \mathcal{L}_p(f^2) &= \sum_{m,m'=0}^{k} \sum_{|I|=m} \sum_{|J|=m'} c_I c_J \mathcal{L}_p(e_I e_J) \\ &= \sum_{m,m'=0}^{k} \sum_{|I|=m} \sum_{|J|=m'} c_I c_J \mathcal{L}_p(e_I)\ell_p(e_J) = \mathcal{L}_p(f)^2 \geq 0. \end{aligned} \tag{4.50}$$

This completes the proof. □

4.8.5 Proof of Lemma 4: irreducible subspaces of $\mathbb{R}^{C_n}$

Proof of Lemma 4. It is easy to see that $\mathrm{Pol}_k(C_n)$ is G_{cube}-invariant for each k. It remains to show that each of these is irreducible. For any $k \in [n]$, let $\epsilon_k \in G_{\mathrm{cube}}$ be

defined by $\epsilon_k = \mathrm{diag}(1, \ldots, -1, \ldots, 1)$ where the -1 is in the k'th position. If $I \subseteq [n]$ we denote by x^I the monomial $\prod_{i \in I} x_i$. Observe that for a given $k \in [n]$ and $I \subseteq [n]$ we have:

$$(\mathrm{id} + \epsilon_k) \cdot x^I = \begin{cases} 2x^I & \text{if } k \notin I \\ 0 & \text{otherwise.} \end{cases}$$

In other words the action of $(\mathrm{id} + \epsilon_k)$ on $\mathbb{R}^{C_n}$ annihilates all monomials involving x_k. Similarly the action of $\prod_{k \in K}(\mathrm{id} + \epsilon_k)$ on $\mathbb{R}^{C_n}$ annihilates all monomials involving any x_k, $k \in K$ since the $(\mathrm{id} + \epsilon_k)$ commute. Now fix some arbitrary k and let V be a (non-zero) invariant subspace of $\mathrm{Pol}_k(C_n)$. We will show that necessarily $V = \mathrm{Pol}_k(C_n)$. Since $V \neq \{0\}$, V contains a nonzero square-free polynomial $p(x)$. By the invariance of V under the permutation action, we can assume that the coefficient of the monomial $x_1 x_2 \cdots x_k$ in p is non-zero, and so p is of the form: $p(x) = c x_1 \cdots x_k + \sum_{|I|=k, I \neq \{1,2,\ldots,k\}} c_I x^I$ with $c \neq 0$. Now note that by the previous observation we have:

$$\left[\prod_{i=k+1}^{n} (\mathrm{id} + \epsilon_i) \right] \cdot p(x) = 2^{n-k} x_1 x_2 \cdots x_k.$$

Hence V is a subspace containing $x_1 x_2 \cdots x_k$ and hence, since V is invariant under the permutation action, it contains every square-free monomial of degree k. It follows that $V = \mathrm{Pol}_k(C_n)$ and so $\mathrm{Pol}_k(C_n)$ is irreducible.

To show the second part of the theorem, we use Proposition 3 from the introduction. Indeed note that $\dim \mathrm{Pol}_k(C_n) = \binom{n}{k}$. Thus Proposition 3 says that if V is an invariant subspace and $\dim V < \binom{n}{k}$ with $k \leq n/2$ then necessarily V is contained in the direct sum

$$\bigoplus_{i=0}^{k-1} (\mathrm{Pol}_i(C_n) \oplus \mathrm{Pol}_{n-i}(C_n)).$$

Thus this means that any $f \in V$ can be decomposed as:

$$f = \sum_{i=0}^{k-1} g_i + g_{n-i}$$

where $g_i \in \mathrm{Pol}_i(C_n)$ and $g_{n-i} \in \mathrm{Pol}_{n-i}(C_n)$. Now note that for $i < n/2$, $\mathrm{Pol}_{n-i}(C_n) \cong e_n(x) \mathrm{Pol}_i(C_n)$ because multiplication by $e_n(x)$ is an involution of $\mathbb{R}^{C_n}$ that sends square-free polynomials of degree i to square-free polynomials of degree $n - i$. Thus we can write $g_{n-i}(x) = e_n(x) h_i(x)$ for some $h_i \in \mathrm{Pol}_i(C_n)$. Thus we get that

$$f(x) = \sum_{i=0}^{k-1} g_i(x) + e_n(x) h_i(x) = g(x) + e_n(x) h(x)$$

where $\deg g \leq k - 1$ and $\deg h \leq k - 1$. □

Chapter 5

Sparse sums of squares and improved semidefinite lifts

In this chapter we show how one can construct improved semidefinite lifts by exhibiting *sparse* sum-of-squares certificates for facet inequalities. We start by showing that the regular N-gon admits an (equivariant) SDP lift of size $O(\log N)$ which matches the lower bound from the previous chapter. In constrast the Lasserre/theta-body hierarchy can be shown to produce a lift of size linear in N.

Motivated by this construction we develop a general framework to produce sparse sum of squares certificates for functions defined on a finite abelian group. We show using this framework that there is an explicit sequence of polytopes in increasing dimensions (so-called trigonometric cyclic polytopes) that admit SDP lifts that are vanishingly smaller than any LP lift. The tools we develop also allow us to prove a conjecture of Laurent from 2003 on the Lasserre hierarchy for maximum cut. This chapter is mostly based on the paper [38], except the first section which is based on part of [36].

5.1 Motivating example: regular polygons

To motivate the study of *sparse* sum-of-squares certificates we first look at the case of regular polygons in $\mathbb{R}^2$ which we considered in the previous chapter. We adopt the same notations considered there and outlined in Section 4.6.2. Recall the facet equation of the regular N-gon (cf. Figure 4-2) expressed in the Fourier basis:

$$\ell = \cos(\pi/N)e_0 - \frac{1}{2}e^{-i\pi/N}e_1 - \frac{1}{2}e^{i\pi/N}e_{-1}.$$

We saw in Theorem 21 that constructing a Rot_N-equivariant SDP lift of the regular N-gon is equivalent to finding a sum-of-squares certificate of ℓ of the form:

$$\ell = \sum_{j=1}^{J} |h_j|^2 \quad \text{where} \quad h_j \in \bigoplus_{k\in K} \mathbb{C}e_k \ \forall j = 1, \ldots, J. \tag{5.1}$$

for some $K \subseteq \mathbb{Z}_N$. The size of the SDP lift is given exactly by $|K|$. In Theorem 22 we proved that any such K must have size at least $\ln(N/2)$. One may wonder whether this bound is tight, i.e., whether there exists a set K of size at most $O(\log N)$ such that (5.1) holds? In this chapter we will show that such a set does indeed exist. In fact the goal of this chapter is to give a combinatorial method to guarantee the existence of sparse sum-of-squares certificates of the form (5.1) in a more general setting.

To give a flavor of how such a sparse certificate may look like we consider the case where N is a power of two, i.e., $N = 2^n$ and consider the problem of finding a set K of logarithmic size in N such that (5.1) holds. In this case one can come up with an explicit such representation which we present in Proposition 8 below. For convenience, define for $k \in \mathbb{Z}_N$, c_k and s_k to be the functions in $\mathbb{C}^{\mathbb{Z}_N}$ that play the role of $\cos(2\pi kt/N)$ and $\sin(2\pi kt/N)$:

$$c_k = \frac{e_k + e_{-k}}{2}, \quad s_k = \frac{e_k - e_{-k}}{2i}.$$

Using these notations the facet equation ℓ can be written as: $\ell = \cos(\pi/N)c_0 - \cos(\pi/N)c_1 - \sin(\pi/N)s_1$.

Proposition 8. *Let ℓ be the facet equation for the regular 2^n-gon, i.e.,*

$$\ell = \cos(\pi/2^n)c_0 - \cos(\pi/2^n)c_1 - \sin(\pi/2^n)s_1,$$

Then we have the following sum-of-squares certificate for ℓ:

$$\ell = \sum_{i=0}^{n-2} \frac{\sin\left(\frac{\pi}{2^n}\right)}{2^i \sin\left(2^{i+1} \cdot \frac{\pi}{2^n}\right)} \left(\cos\left(\frac{\pi}{2^{n-i}}\right) c_0 - \cos\left(\frac{\pi}{2^{n-i}}\right) c_{2^i} - \sin\left(\frac{\pi}{2^{n-i}}\right) s_{2^i}\right)^2. \quad (5.2)$$

Proof. The identity (5.2) can be proved using induction. We will omit the proof since later we will give a more general result concerning general N-gons (and not just the case where N is a power of two). □

What is important to note in (5.2) is that only the elements c_0, c_{2^i} and s_{2^i} for $i = 0, \ldots, n-2$ are used in the sum-of-squares certificate. In other words the sparsity pattern of the certificate is $K = \{0\} \cup \{\pm 2^i, i = 0, \ldots, n-2\}$ which has size $|K| = 2n-1$. Also that the certificate (5.2) has high-degree (it uses monomials of degree $N/4 = 2^{n-2}$) despite being very sparse. We have actually shown in the previous chapter (see Remark 10) that any certificate of the form (5.1) for the regular N-gon must have degree at least $N/4$.

Lifts Before concluding this example we give the explicit SDP lift obtained from the sum-of-squares certificate (5.2). Recall from Chapter 2, Section 2.2.5 how the SDP lift is constructed from the sum-of-squares certificates via the notion of *pseudo-expectation* (Equation (2.20)). If we let $V_i = \mathbb{C}e_0 \oplus \mathbb{C}e_{2^i} \oplus \mathbb{C}e_{-2^i}$ then our lift has the following abstract form:

$$\begin{aligned} \mathrm{conv}(\mathcal{X}_{2^n}) = \big\{(\widetilde{E}(c_1), \widetilde{E}(s_1)) \ : \widetilde{E} \in \mathbb{C}^{2^n} \text{ with } \widetilde{E}(e_0) = 1, \\ \widetilde{E}(f^2) \geq 0 \ \forall f \in V_i, i = 0, \ldots, n-2\big\}. \end{aligned} \quad (5.3)$$

The vector $\widetilde{E}$ is indexed by the elements $(e_k)_{k \in \mathbb{Z}_{2^n}}$, and one should think of $\widetilde{E}(e_k)$ as the (pseudo)-expectation of the function $e_k(t)$ on $t \in \mathbb{Z}_{2^n}$. From (5.3) the vector $\widetilde{E}$ has size 2^n, however as we will see only a small number of components (linear in n) will actually be needed to express the final lift. To make the notations lighter we will denote $m_k := \widetilde{E}(e_k) \in \mathbb{C}$ for any $k \in \mathbb{Z}_N$. Note that $m_{-k} = \overline{m_k}$. We now need to express the constraint that $\widetilde{E}(f^2) \geq 0 \ \forall f \in V_i$ using the (m_k). Using the basis (e_0, e_{2^i}, e_{-2^i}) of V_i we see that the matrix of the quadratic form $f \in V_i \mapsto \widetilde{E}(f^2)$ is given by:

$$\begin{bmatrix} m_0 & m_{2^i} & \overline{m_{2^i}} \\ \overline{m_{2^i}} & m_0 & \overline{m_{2^{i+1}}} \\ m_{2^i} & m_{2^{i+1}} & m_0 \end{bmatrix}.$$

To see how this matrix is constructed note for example that the $(2,3)$ entry is given by $\widetilde{E}(\overline{e_{2^i}} e_{-2^i}) = \widetilde{E}(e_{-2^{i+1}}) = \overline{\widetilde{E}(e_{2^{i+1}})} = \overline{m_{2^{i+1}}}$. The lift (5.3) of the regular 2^n-gon

thus can be written as (where y_i plays the role of m_{2^i}):

$$\operatorname{conv}(\mathcal{X}_{2^n}) = \Bigg\{ (\operatorname{Re}[y_0], \operatorname{Im}[y_0]) : \exists y_1, \ldots, y_{n-2} \in \mathbb{C}, y_{n-1} \in \mathbb{R} \text{ such that } \begin{bmatrix} 1 & y_{i-1} & \overline{y_{i-1}} \\ \overline{y_{i-1}} & 1 & \overline{y_i} \\ y_{i-1} & y_i & 1 \end{bmatrix} \in \mathbf{H}^3_+ \quad \text{for } i = 1, 2, \ldots, n-2 \text{ and } \begin{bmatrix} 1 & y_{n-2} & \overline{y_{n-2}} \\ \overline{y_{n-2}} & 1 & y_{n-1} \\ y_{n-2} & y_{n-1} & 1 \end{bmatrix} \in \mathbf{H}^3_+ \Bigg\}. \tag{5.4}$$

Remark 12 (Real equivariant SDP lift). Observe that Proposition 8 actually gives a *real* sum-of-squares certificate of ℓ, i.e., the functions h_j in $\ell = \sum_j |h_j|^2$ are real-valued. This sum-of-squares certificate can be used to get a SDP lift of the regular 2^n-gon using the cone of real symmetric psd matrices (instead of Hermitian psd matrices). The real SDP lift can be shown to take the form ($\mathbf{S}^3_+$ denotes the cone of 3×3 real symmetric positive semidefinite matrices):

$$\operatorname{conv}(\mathcal{X}_{2^n}) = \Bigg\{ (x_0, y_0) : \exists (x_i, y_i)_{i=1}^{n-2}, x_{n-1} \in \mathbb{R}, \begin{bmatrix} 1 & x_{i-1} & y_{i-1} \\ x_{i-1} & \frac{1+x_i}{2} & \frac{y_i}{2} \\ y_{i-1} & \frac{y_i}{2} & \frac{1-x_i}{2} \end{bmatrix} \in \mathbf{S}^3_+ \quad \text{for } i = 1, 2, \ldots, n-2 \text{ and } \begin{bmatrix} 1 & x_{n-2} & y_{n-2} \\ x_{n-2} & \frac{1+x_{n-1}}{2} & 0 \\ y_{n-2} & 0 & \frac{1+x_{n-1}}{2} \end{bmatrix} \in \mathbf{S}^3_+ \Bigg\}. \tag{5.5}$$

Remark 13. One can also get from Proposition 8 a slightly different SDP lift of the regular 2^n-gon that involves only a single LMI of size $2n-1$, rather than $n-1$ LMIs of size 3 each. If we let $K = \{0\} \cup \{\pm 2^i, i = 0, \ldots, n-2\}$ (which is the sparsity pattern of the sum-of-squares certificate of Proposition 8) then this lift takes the form:

$$\operatorname{conv}(\mathcal{X}_{2^n}) = \left\{ (\operatorname{Re}[m_1], \operatorname{Im}[m_1]) : m_0 = 1 \text{ and } [m_{k'-k}]_{k,k' \in K} \in \mathbf{H}^{2n-1}_+ \right\}.$$

For example for $N = 16$ we get that $\operatorname{conv}(\mathcal{X}_{16})$ is the set of $(\operatorname{Re}[m_1], \operatorname{Im}[m_1]) \in \mathbb{R}^2$

such that the following 7×7 Hermitian matrix is positive semidefinite:

$$\begin{bmatrix} 1 & m_1 & \overline{m_1} & m_2 & \overline{m_2} & m_4 & \overline{m_4} \\ \overline{m_1} & 1 & \overline{m_2} & m_1 & \overline{m_3} & m_3 & \overline{m_5} \\ m_1 & m_2 & 1 & m_3 & \overline{m_1} & m_5 & \overline{m_3} \\ \overline{m_2} & \overline{m_1} & \overline{m_3} & 1 & \overline{m_4} & m_2 & \overline{m_6} \\ m_2 & m_3 & m_1 & m_4 & 1 & m_6 & \overline{m_2} \\ \overline{m_4} & \overline{m_3} & \overline{m_5} & \overline{m_2} & \overline{m_6} & 1 & m_8 \\ m_4 & m_5 & m_3 & m_6 & m_2 & m_8 & 1 \end{bmatrix}.$$

Note that the auxiliary variables are $m_2, m_3, m_4, m_5, m_6, m_8$ and that m_8 is a real variable whereas the others are complex.

5.2 The setting of finite abelian groups

We now consider a general setting where one is interested in finding sparse sum-of-squares certificates for sparse nonnegative functions, and we describe a graph-theoretic method to find such certificates. This will allow us to construct small (equivariant) SDP lifts for certain classes of moment polytopes that we describe in detail later.

Let G be a finite abelian group. It is well-known that any function $f : G \to \mathbb{C}$ admits a Fourier decomposition where the Fourier basis consists of the *characters* of G. Such a decomposition takes the form

$$f(x) = \sum_{\chi \in \widehat{G}} \widehat{f}(\chi)\chi(x) \quad \forall x \in G$$

where $\widehat{G}$ is the set of characters of G (known as the *dual group* of G) and $\widehat{f}(\chi)$ are the Fourier coefficients of f. The function $f : G \to \mathbb{C}$ is called *Fourier-sparse* if only a few of its Fourier coefficients are nonzero. More precisely we say that f has Fourier support $\mathcal{S} \subseteq \widehat{G}$ if $\widehat{f}(\chi) = 0$ whenever $\chi \notin \mathcal{S}$.

In this chapter we are concerned with functions $f : G \to \mathbb{C}$ that are *Fourier-sparse* and *nonnegative*, i.e., $f(x) \geq 0$ for all $x \in G$. If f is a nonnegative function on G, a *sum-of-squares certificate* for the nonnegativity of f has the form

$$f(x) = \sum_{j=1}^{J} |f_j(x)|^2 \quad \forall x \in G \tag{5.6}$$

where $f_j : G \to \mathbb{C}$. When the function f is Fourier-sparse, it is natural to ask whether f admits a sum-of-squares certificate that is also Fourier-sparse, i.e., where all the functions f_j have Fourier support on a common "small" set $\mathcal{T} \subseteq \widehat{G}$. This leads us to

the main problem of interest in this chapter.

> Given $\mathcal{S} \subseteq \widehat{G}$, find a subset $\mathcal{T} \subseteq \widehat{G}$ such that any nonnegative function $G \to \mathbb{R}_+$ with Fourier support $\mathcal{S}$ admits a sum-of-squares certificate with Fourier support $\mathcal{T}$. (P)

Our main result is to give a sufficient condition for a set $\mathcal{T}$ to satisfy the requirement above for a given $\mathcal{S}$. The condition is expressed in terms of *chordal covers* of the Cayley graph $\mathrm{Cay}(\widehat{G}, \mathcal{S})$. Recall that the Cayley graph $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ is the graph where nodes correspond to elements of $\widehat{G}$ and where χ, χ' are connected by an edge if $\chi^{-1}\chi' \in \mathcal{S}$. The following is our main result.

Theorem 25. *Let $\mathcal{S} \subseteq \widehat{G}$, let Γ be a chordal cover of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$, and for each maximal clique $\mathcal{C}$ of Γ, let $\chi_{\mathcal{C}}$ be an element of $\widehat{G}$. Define*

$$\mathcal{T} := \mathcal{T}(\Gamma, \{\chi_{\mathcal{C}}\}) = \bigcup_{\mathcal{C}} \chi_{\mathcal{C}}\mathcal{C} \tag{5.7}$$

where the union is over all the maximal cliques of Γ and where $\chi_{\mathcal{C}}\mathcal{C} := \{\chi_{\mathcal{C}}\chi : \chi \in \mathcal{C}\}$ is the translation of $\mathcal{C}$ by $\chi_{\mathcal{C}}$. Then any nonnegative function with Fourier support $\mathcal{S}$ admits a sum-of-squares certificate with Fourier support $\mathcal{T}$.

Theorem 25 gives a way to construct a set $\mathcal{T}$ that satisfies the condition in (P) for a given $\mathcal{S} \subseteq \widehat{G}$. Such a construction proceeds in two steps: first choose a chordal cover Γ of the graph $\mathrm{Cay}(\widehat{G}, \mathcal{S})$, and then choose elements $\chi_{\mathcal{C}} \in \widehat{G}$ for each maximal clique $\mathcal{C}$ of $\widehat{G}$. Different choices of Γ and $\{\chi_{\mathcal{C}}\}$ will in general lead to different sets $\mathcal{T}(\Gamma, \{\chi_{\mathcal{C}}\})$. When using Theorem 25, one wants to find a good choice of Γ and $\{\chi_{\mathcal{C}}\}$ such that the resulting set $\mathcal{T}(\Gamma, \{\chi_{\mathcal{C}}\})$ is as small as possible (or has other desirable properties).

One of the main strengths of Theorem 25 is in the ability to choose the elements $\{\chi_{\mathcal{C}}\}$. In fact the conclusion of Theorem 25 is almost trivial if $\chi_{\mathcal{C}} = 1_{\widehat{G}}$ for all $\mathcal{C}$, since in this case it simply says that any nonnegative function has a sum-of-squares certificate supported on $\widehat{G}$, which is easy to see since G is finite. As we will see in the applications, it is the ability to *translate* the cliques $\mathcal{C}$ of Γ via the choice of $\chi_{\mathcal{C}}$ that is key in Theorem 25 and allows us to obtain interesting results. Equation (5.7) gives us the intuition behind a good choice of $\{\chi_{\mathcal{C}}\}$: in order to minimize the cardinality of $\mathcal{T}(\Gamma, \{\chi_{\mathcal{C}}\})$ one would like to find the translations $\chi_{\mathcal{C}}$ that maximize the total overlap of the cliques (i.e., minimize the cardinality of their union).

Before describing the main idea behind Theorem 25 and its proof, we illustrate how one can use Theorem 25 in two important special cases, namely $G = \mathbb{Z}_2^n$ (the boolean hypercube) and $G = \mathbb{Z}_N$.

- **Boolean hypercube**: Consider the case $G = \{-1, 1\}^n \cong \mathbb{Z}_2^n$. The Fourier expansion of functions on $\{-1, 1\}^n$ takes the form

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{i \in S} x_i. \tag{5.8}$$

A function f is said to have degree d if $\widehat{f}(S) = 0$ for all S such that $|S| > d$. Many combinatorial optimization problems correspond to optimizing a certain function f over $\{-1, 1\}^n$. For example the maximum cut problem in graph theory consists in optimizing a *quadratic function* over $\{-1, 1\}^n$. In [72] Laurent conjectured that any nonnegative quadratic function on the hypercube is a sum of squares of functions of degree at most $\lceil n/2 \rceil$. Using our notations, this corresponds to asking whether for $\mathcal{S} = \{S \subseteq [n] : |S| = 0 \text{ or } 2\}$ one can find $\mathcal{T} \subseteq \{S \subseteq [n] : |S| \leq \lceil n/2 \rceil\}$ such that the conclusion of Theorem 25 holds. By studying chordal covers of the Cayley graph $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ we are able to answer this question positively.

Theorem 26. *Any nonnegative quadratic function on $\{-1, 1\}^n$ is a sum-of-squares of polynomials of degree at most $\lceil n/2 \rceil$.*

Note that Blekherman et al. [9] previously showed a weaker version of the conjecture that allows for multipliers: They showed that for any nonnegative quadratic function f on the hypercube, there exists h sum-of-squares such that $h(x)f(x)$ is a sum of squares of polynomials of degree at most $\lceil n/2 \rceil$.

- **Trigonometric polynomials**: Another important application that we consider is the case where $G = \mathbb{Z}_N$, the (additive) group of integers modulo N. The Fourier decomposition of a function $f : \mathbb{Z}_N \to \mathbb{C}$ is the usual discrete Fourier transform and takes the form

$$f(x) = \sum_{k \in \mathbb{Z}_N} \widehat{f}(k) e^{2i\pi kx/N} \tag{5.9}$$

where $\widehat{f}(k)$ are the Fourier coefficients of f. Nonnegative trigonometric polynomials play an important role in many areas such as in signal processing [31], but also in convex geometry [102, 3], in their relation to (trigonometric) cyclic polytopes. We are interested in nonnegative functions on $G = \mathbb{Z}_N$ of degree at most d, i.e., functions with Fourier support $\mathcal{S} = \{-d, -(d-1), \ldots, d-1, d\}$. By studying chordal covers of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ (which is nothing but the dth power of the cycle graph) and using Theorem 25 we are able to establish the following result.

Theorem 27. *Let N and d be two integers and assume that d divides N. Then there exists $\mathcal{T} \subseteq \mathbb{Z}_N$ with $|\mathcal{T}| \leq 3d \log(N/d)$ such that any nonnegative function on $\mathbb{Z}_N$ of degree at most d has a sum-of-squares certificate with Fourier support $\mathcal{T}$.*

As will be clear later when we describe the dual point of view, the case of the regular N-gon corresponds to the choice $d = 1$.

Remark 14. Note that if one is interested in functions of degree at most d on $\mathbb{Z}_N$ and d does not divide N, then one can still apply Theorem 27 with d' instead of d, where d' is the smallest divisor of N that is greater than d.

Dual point of view and moment polytopes Theorem 25 can be interpreted from the dual point of view as giving a semidefinite programming description of certain moment polytopes. If $\mathcal{S} \subseteq \widehat{G}$, define the *moment polytope* $\mathcal{M}(G, \mathcal{S})$ to be the set of $\mathcal{S}$-moments of probability distributions on G, i.e.,

$$\mathcal{M}(G, \mathcal{S}) = \Big\{ \big(\mathbb{E}_{x \sim \mu}[\chi(x)]\big)_{\chi \in \mathcal{S}} \in \mathbb{C}^{\mathcal{S}} : \mu \text{ a probability measure supported on } G \Big\}.$$

Note that $\mathcal{M}(G, \mathcal{S})$ is a polytope since it can be equivalently expressed as

$$\mathcal{M}(G, \mathcal{S}) = \operatorname{conv}\Big\{ (\chi(x))_{\chi \in \mathcal{S}} \in \mathbb{C}^{\mathcal{S}} : x \in G \Big\}. \tag{5.10}$$

For example if $G = \mathbb{Z}_N$ and $\mathcal{S} = \{-1, 1\}$, then the moment polytope $\mathcal{M}(G, \mathcal{S})$ is nothing but the regular N-gon (up to a linear transformation). From a geometric point of view, nonnegative functions $f : G \to \mathbb{R}_+$ with Fourier support $\mathcal{S}$ correspond to valid linear inequalities for the polytope $\mathcal{M}(G, \mathcal{S})$. By giving a sum-of-squares characterization for all valid inequalities of $\mathcal{M}(G, \mathcal{S})$ Theorem 25 allows us to obtain a semidefinite programming description of $\mathcal{M}(G, \mathcal{S})$. The following statement can be obtained from Theorem 25 by duality. (We call this result "Theorem 25D" to reflect that it is a dual version of "Theorem 25" and adopt this numbering convention throughout the chapter.)

Theorem 25D. *Let $\mathcal{S} \subseteq \widehat{G}$ and let $\mathcal{T} = \mathcal{T}(\Gamma, \{\chi_{\mathcal{C}}\})$ be as defined in Theorem 25. Then we have the following semidefinite programming description of the moment polytope $\mathcal{M}(G, \mathcal{S})$:*

$$\mathcal{M}(G, \mathcal{S}) = \Big\{ (\ell_\chi)_{\chi \in \mathcal{S}} \;:\; \exists (y_\chi)_{\chi \in \mathcal{T}^{-1}\mathcal{T}} \text{ such that } y_\chi = \ell_\chi \text{ for all } \chi \in \mathcal{S}\ , \text{ and } y_{1_{\widehat{G}}} = 1, \text{ and } \big[y_{\overline{\chi}\chi'}\big]_{\chi, \chi' \in \mathcal{T}} \succeq 0 \Big\}. \tag{5.11}$$

In terms of positive semidefinite lifts, Equation (5.11) shows that $\mathcal{M}(G, \mathcal{S})$ has a *Hermitian positive semidefinite lift* of size $|\mathcal{T}|$. We now illustrate this dual point of view for the two applications mentioned above, $G = \{-1, 1\}^n$ and $G = \mathbb{Z}_N$.

- For the case of the boolean hypercube $G = \{-1, 1\}^n$, if $\mathcal{S} = \{S \subseteq [n] : |S| = 0 \text{ or } 2\}$, the moment polytope $\mathcal{M}(\{-1, 1\}^n, \mathcal{S} \setminus \{\emptyset\})$ is nothing but the *cut polytope* for the complete graph on n vertices. We use the notation CUT_n for this polytope, i.e.,

  $$\mathrm{CUT}_n = \operatorname{conv}\Big\{ (x_i x_j)_{i<j} \in \mathbb{R}^{\binom{n}{2}} : x \in \{-1, 1\}^n \Big\}.$$

  From the dual point of view, Theorem 26 shows that the $\lceil n/2 \rceil$ level of the Lasserre hierarchy for the cut polytope is exact. This bound is tight since Laurent showed in [72] that at least $\lceil n/2 \rceil$ levels are needed.

  Theorem 26D. *The $\lceil n/2 \rceil$ level of the Lasserre hierarchy for the cut polytope CUT_n is exact.*

- Consider now the case $G = \mathbb{Z}_N$ and $\mathcal{S} = \{-d, -(d-1), \ldots, d-1, d\}$. Here the moment polytope $\mathcal{M}(G, \mathcal{S})$ is the *trigonometric cyclic polytope* of degree d. We use the notation $TC(N, 2d)$ for this polytope, i.e.,

$$TC(N, 2d) = \operatorname{conv}\Big\{TM(2\pi x/N) : x = 0, 1, \ldots, N-1\Big\} \subset \mathbb{R}^{2d}, \tag{5.12}$$

where $TM(\theta)$ is the degree d trigonometric moment curve

$$TM(\theta) = \Big(\cos(\theta), \sin(\theta), \cos(2\theta), \sin(2\theta), \ldots, \cos(d\theta), \sin(d\theta)\Big).$$

When interpreted from the dual point of view, Theorem 27 shows that $TC(N, 2d)$ has a Hermitian positive semidefinite lift of size at most $3d \log(N/d)$.

Theorem 27D. *Let N and d be two integers and assume that d divides N. The trigonometric cyclic polytope $TC(N, 2d)$ defined in (5.12) has a Hermitian positive semidefinite lift of size at most $3d \log(N/d)$.*

Note that in the case $d = 1$ the polytope $TC(N, 2d)$ is nothing but the regular N-gon in the plane. Theorem 27D thus recovers, and extends to the case where N is not a power of two, the result from Section 5.1 giving a SDP lift of the regular N-gon of size $O(\log N)$ for $N = 2^n$.

For $d > 1$ our result is, as far as we are aware, the first nontrivial semidefinite programming lift of a cyclic polytope. Furthermore, in the regime where $N = d^2$ our lift is provably smaller than any linear programming lift. Indeed, since $TC(d^2, 2d)$ is d-neighborly [45], a lower bound from [40] concerning neighborly polytopes shows that any linear programming lift of $TC(d^2, 2d)$ must have size at least $\Omega(d^2)$, whereas our semidefinite programming lift in this case has size $O(d \log d) = o(d^2)$. To the best of our knowledge this gives the first example of a family of polytopes $(P_d)_{d \in \mathbb{N}}$ in increasing dimensions where $\mathrm{xc}_{\mathrm{SDP}}(P_d) = o(\mathrm{xc}_{\mathrm{LP}}(P_d))$ where $\mathrm{xc}_{\mathrm{SDP}}$ and $\mathrm{xc}_{\mathrm{LP}}$ are respectively the SDP and LP extension complexity.

Corollary 1. *There exists an explicit family $(P_d)_{d \in \mathbb{N}}$ of polytopes where $P_d \subset \mathbb{R}^{2d}$ such that*

$$\frac{\mathrm{xc}_{SDP}(P_d)}{\mathrm{xc}_{LP}(P_d)} = O\left(\frac{\log d}{d}\right).$$

The only nontrivial linear programming lift for cyclic polytopes that we are aware of is a construction by Bogomolov et al. [12] for the polytope

$$\operatorname{conv}\{(i, i^2, \ldots, i^d) : i = 1, \ldots, N\}$$

which has size $(\log N)^{\lfloor d/2 \rfloor}$.

Main ideas We now briefly describe the main ideas behind Theorem 25, which can be summarized in three steps.

1. *A sum-of-squares certificate with a sparse Gram matrix*: Given a nonnegative function $f : G \to \mathbb{R}_+$ it is easy to see, since G is finite, that f can be written as a sum-of-squares. When the function f has Fourier support $\mathcal{S}$, one can show that f admits a specific sum-of-squares representation where the *Gram matrix* Q, in the basis of characters, is sparse according to the graph $\mathrm{Cay}(\widehat{G}, \mathcal{S})$.

2. *Chordal completion*: Let Γ be a chordal cover of the graph $\mathrm{Cay}(\widehat{G}, \mathcal{S})$. Using well-known results concerning positive semidefinite matrices that are sparse according to a chordal graph [55, 56] (see Section 5.3.2 for more details) one can decompose the Gram matrix Q into a sum of positive-semidefinite matrices, where each matrix is supported on a maximal clique of Γ. In terms of sum-of-squares representation, this means that the function f can be written as

$$f = \sum_j |f_j|^2 \tag{5.13}$$

   where each f_j has Fourier support on a maximal clique $\mathcal{C}_j \subset \widehat{G}$ of Γ.

3. *Translation of cliques*: The problem with the decomposition (5.13) is that even though each maximal clique $\mathcal{C}_j$ might be small, the union of the $\mathcal{C}_j$'s might be large, and thus the total Fourier support of (5.13) might be large (in fact the union of the $\mathcal{C}_j$'s is the whole $\widehat{G}$). In order to reduce the total Fourier support of the sum-of-squares certificate (5.13), we use the following simple but crucial observation: if h is a function with Fourier support $\mathcal{C}$ and if $\chi \in \widehat{G}$ then χh has Fourier support $\chi\mathcal{C}$ and we have $|\chi h|^2 = |h|^2$. Thus if for each maximal clique $\mathcal{C}_j$ of Γ we choose a certain $\chi_j \in \widehat{G}$ then, by translating each term in (5.13) by χ_j we obtain a sum-of-squares representation of f of the form $f = \sum_j |\widetilde{h}_j|^2$ where $\widetilde{h}_j$ has Fourier support $\chi_j\mathcal{C}_j$. Having chosen the χ_j such that $\chi_j\mathcal{C}_j \subseteq \mathcal{T}$ for all maximal cliques $\mathcal{C}_j$ (cf. Equation (5.7)), we get a representation of f as a sum-of-squares of functions with Fourier support $\mathcal{T}$.

Organization The rest of the chapter is organized as follows. Section 5.3 starts by giving a brief review of Fourier analysis of finite abelian groups, as well as a review of chordal graphs, chordal covers and the main results concerning decomposition/matrix completion with chordal sparsity structure [55, 56]. In Section 5.4 we prove our main result, Theorem 25. In Section 5.5 we look at the case of the hypercube $G = \{-1, 1\}^n$ mentioned earlier, and we look in particular at quadratic functions on the hypercube. We give an explicit chordal cover for the corresponding Cayley graph and we show how it leads to a proof of Laurent's conjecture. In Section 5.6 we look at the special case $G = \mathbb{Z}_N$ and functions of degree d. We give an explicit chordal cover for the corresponding graphs, and we discuss the consequences concerning positive semidefinite lifts of the trigonometric cyclic polytope.

5.3 Background: Fourier analysis and chordal completion

In this section we present some background material needed for this chapter: we first recall some of the basic results and terminology concerning Fourier analysis on finite abelian groups [90, 97], then we review the definition of chordal graph and the main results concerning sparse positive semidefinite matrices and matrix completion.

5.3.1 Fourier analysis on finite abelian groups

Let G be a finite abelian group which we denote multiplicatively, and let $\mathbb{C}^G$ be the vector space of complex-valued functions on G. A *character* χ of G is a *group homomorphism* $\chi : G \to (\mathbb{C}^*, \times)$, i.e., it is an element of $\mathbb{C}^G$ which satisfies:

$$\chi(xy) = \chi(x)\chi(y) \quad \forall x, y \in G.$$

Since G is abelian, one can easily show that the (pointwise) product of two characters is a character and that the (pointwise) inverse of a character is again a character. Thus if we denote by $\widehat{G}$ the set of characters of G, then $\widehat{G}$ forms an abelian group, where the group operation corresponds to pointwise multiplication. The group $\widehat{G}$ is known as the *dual group* of G. Observe that since G is finite, if χ is a character then for any $x \in G$ we have $\chi(x)^{|G|} = \chi(x^{|G|}) = \chi(1_G) = 1$, which implies that $|\chi(x)| = 1$. It follows that the inverse of a character χ is simply its (pointwise) complex conjugate $\overline{\chi}$.

Consider the standard inner product on $\mathbb{C}^G$ defined by

$$\langle f, g \rangle = \frac{1}{|G|} \sum_{x \in G} f(x)\overline{g(x)} \quad \forall f, g \in \mathbb{C}^G. \tag{5.14}$$

A crucial property of the set of characters $\widehat{G}$ is that they form an orthonormal basis of $\mathbb{C}^G$, which is called the *Fourier basis* of G. Note that this implies in particular that $|\widehat{G}| = |G|$. We summarize this in the following theorem.

Theorem 28. *Let G be a finite abelian group and let $\widehat{G}$ be the set of characters of G. Then $\widehat{G}$ is an abelian group with pointwise multiplication. Furthermore $|\widehat{G}| = |G|$ and $\widehat{G}$ forms an orthonormal basis of $\mathbb{C}^G$ for the standard inner product (5.14).*

We now illustrate the previous theorem in the two examples $G = \{-1, 1\}^n$ (the hypercube) and $G = \mathbb{Z}_N$ presented in the introduction.

Example 11 (Fourier analysis on the hypercube). Let $G = \{-1, 1\}^n$ be the hypercube in dimension n which forms a group of size 2^n under componentwise multiplication, isomorphic to $\mathbb{Z}_2^n$. Observe that if S is a subset of $[n]$ then the function χ_S defined by

$$\chi_S : \{-1, 1\}^n \to \mathbb{C}^*, \quad \chi_S(x) = \prod_{i \in S} x_i$$

satisfies $\chi_S(xy) = \chi_S(x)\chi_S(y)$, and thus is a character of G. For example $\chi_\emptyset$ is the constant function equal to 1, and $\chi_{[n]}$ is the function $\chi_{[n]}(x) = x_1 \dots x_n$. One can show that these are all the characters of G, i.e., $\widehat{G} = \{\chi_S, S \subseteq [n]\}$. Thus the decomposition of a function $f : \{-1,1\}^n \to \mathbb{C}$ in the basis of characters takes the form

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{i \in S} x_i,$$

where $\widehat{f}(S)$ are the Fourier coefficients of f. $\lozenge$

Example 12 (Fourier analysis on $\mathbb{Z}_N$). Let N be an integer and consider the (additive) group $G = \mathbb{Z}_N$ of integers modulo N. For $k \in \mathbb{Z}_N$, define χ_k by

$$\chi_k : \mathbb{Z}_N \to \mathbb{C}^*, \quad \chi_k(x) = e^{2i\pi kx/N}.$$

Note that χ_k satisfies $\chi_k(x+y) = \chi_k(x)\chi_k(y)$ and thus χ_k is a character of $\mathbb{Z}_N$. It is not hard to show that any character χ of $\mathbb{Z}_N$ actually must have the form $\chi = \chi_k$ for some $k \in \mathbb{Z}_N$. Thus the dual group $\widehat{\mathbb{Z}_N}$ of $\mathbb{Z}_N$ is $\widehat{\mathbb{Z}_N} = \{\chi_k, k \in \mathbb{Z}_N\}$. Note that $\chi_k\chi_{k'} = \chi_{k+k'}$ and $(\chi_k)^{-1} = \overline{\chi_k} = \chi_{-k}$, and thus $\widehat{\mathbb{Z}_N}$ is isomorphic to $\mathbb{Z}_N$. According to Theorem 28, any function $f : \mathbb{Z}_N \to \mathbb{C}$ can be decomposed in the basis of characters

$$f(x) = \sum_{k \in \mathbb{Z}_N} \widehat{f}(k) e^{2i\pi kx/N} \quad \forall x \in \mathbb{Z}_N.$$

This decomposition is nothing but the well-known Fourier decomposition of discrete signals of length N. $\lozenge$

For a general finite abelian group G, the Fourier decomposition of a function $f : G \to \mathbb{C}$, in the orthonormal basis of characters takes the form

$$f(x) = \sum_{\chi \in \widehat{G}} \widehat{f}(\chi)\chi(x).$$

The coefficients $\widehat{f}(\chi)$ are the *Fourier coefficients* of f. By orthonormality of the basis of characters, we have for any $\chi \in \widehat{G}$,

$$\widehat{f}(\chi) = \langle f, \chi \rangle = \frac{1}{|G|} \sum_{x \in G} f(x)\overline{\chi(x)}.$$

The *Fourier support* of a function f, denoted $\operatorname{supp} \widehat{f}$ is the set of characters χ for which $\widehat{f}(\chi) \neq 0$, i.e.,

$$\operatorname{supp} \widehat{f} = \{\chi \in \widehat{G} : \widehat{f}(\chi) \neq 0\}.$$

5.3.2 Chordal graphs and matrix completion

In this section we recall some of the main results concerning sparse matrix decomposition and matrix completion with a chordal sparsity structure. For more details, we

refer the reader to [56, 55] and [1].

Chordal graphs Let $\mathcal{G} = (V, E)$ be a graph. The graph $\mathcal{G}$ is called *chordal* if any cycle of length at least four has a chord. A *chordal cover* (also called *triangulation*) of $\mathcal{G}$ is a graph $\mathcal{G}' = (V, E')$ where $E \subset E'$ and where $\mathcal{G}'$ is chordal. Figure 5-1 shows a non-chordal graph $\mathcal{G}$ on four vertices and a chordal cover $\mathcal{G}'$ of $\mathcal{G}$.

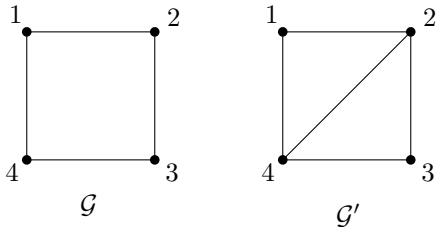

Figure 5-1: A non-chordal graph $\mathcal{G}$ and a chordal cover $\mathcal{G}'$ of $\mathcal{G}$.

A subset $\mathcal{C} \subseteq V$ is a *clique* in $\mathcal{G}$ if $\{i, j\} \in E$ for all $i, j \in \mathcal{C}$, $i \neq j$. The clique $\mathcal{C}$ is called *maximal* if it is not a strict subset of another clique $\mathcal{C}'$ of $\mathcal{G}$. For example the maximal cliques of the graph $\mathcal{G}'$ shown in Figure 5-1 are $\{1, 2, 4\}$ and $\{2, 3, 4\}$.

Sparse matrices Let $Q \in \mathbf{H}_+^V$ be a Hermitian positive semidefinite matrix where rows and columns are indexed by some set V. Assume furthermore that Q is sparse according to some graph $\mathcal{G} = (V, E)$, i.e.,

$$Q_{ij} \neq 0, i \neq j \Rightarrow \{i, j\} \in E.$$

One of the main tools that we use in the proof of our main theorem is a result from [55, 56] which allows us to decompose sparse positive semidefinite matrices as a sum of positive semidefinite matrices supported on a small subset of rows/columns. We say that a Hermitian matrix A is supported on $\mathcal{C} \subseteq V$ if $A_{ij} = 0$ whenever $i \notin \mathcal{C}$ or $j \notin \mathcal{C}$. The result can be stated as follows.

Theorem 29. *([55, 56]) Let Q be a Hermitian positive semidefinite matrix, and assume that Q is sparse according to some graph $\mathcal{G}$. Assume furthermore that $\mathcal{G}$ is chordal. Then for every maximal clique $\mathcal{C}$ of $\mathcal{G}$ there exists a Hermitian positive semidefinite matrix $Q_{\mathcal{C}}$ supported on $\mathcal{C}$ such that*

$$Q = \sum_{\mathcal{C}} Q_{\mathcal{C}}. \tag{5.15}$$

Remark 15. If the sparsity pattern $\mathcal{G}$ of Q is not chordal, one can still apply the previous theorem by considering a chordal cover $\mathcal{G}'$ of $\mathcal{G}$. Indeed if Q is sparse according to $\mathcal{G}$ then it also clearly sparse according to $\mathcal{G}'$, since $\mathcal{G} \subseteq \mathcal{G}'$. In this case the summation (5.15) is over the maximal cliques of $\mathcal{G}'$.

Example 13. We can illustrate the previous theorem with a simple 4×4 matrix. Let Q be the 4×4 Hermitian positive semidefinite matrix given by

$$Q = \begin{bmatrix} 2 & 1-i & 0 & 1+i \\ 1+i & 2 & 1-i & 0 \\ 0 & 1+i & 2 & 1-i \\ 1-i & 0 & 1+i & 2 \end{bmatrix}.$$

Note that Q is sparse according to the "square graph" $\mathcal{G}$ shown in Figure 5-1(left). Since $\mathcal{G}$ is not chordal we cannot directly apply Theorem 29 with $\mathcal{G}$, but we can apply it with $\mathcal{G}'$ shown in Figure 5-1(right) which is a chordal cover of $\mathcal{G}$. In this case Theorem 29 asserts that one can decompose Q as a sum of two positive semidefinite matrices supported respectively on the maximal cliques, $\{1, 2, 4\}$ and $\{2, 3, 4\}$. For this example, it is not hard to find an explicit decomposition, for example we can verify that

$$Q = \underbrace{\begin{bmatrix} 2 & 1-i & 0 & 1+i \\ 1+i & 1 & 0 & i \\ 0 & 0 & 0 & 0 \\ 1-i & -i & 0 & 1 \end{bmatrix}}_{\succeq 0} + \underbrace{\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 1-i & -i \\ 0 & 1+i & 2 & 1-i \\ 0 & i & 1+i & 1 \end{bmatrix}}_{\succeq 0}.$$

$\Diamond$

Matrix completion One can also state Theorem 29 in its dual form, in terms of the *matrix completion problem*. Given a graph $\mathcal{G} = (V, E)$, a *$\mathcal{G}$-partial matrix* X is a matrix where only the diagonal entries, as well as the entries X_{ij} for $\{i, j\} \in E$ are specified. Given a $\mathcal{G}$-partial matrix X, the positive semidefinite matrix completion problem asks whether X can be completed into a full $|V| \times |V|$ Hermitian matrix that is positive semidefinite. Clearly a necessary condition for such a completion to exist is that $X[\mathcal{C}, \mathcal{C}] \succeq 0$ for all cliques $\mathcal{C}$ of $\mathcal{G}$ (note that if $\mathcal{C}$ is a clique of $\mathcal{G}$, then all the entries of $X[\mathcal{C}, \mathcal{C}]$ are specified). When $\mathcal{G}$ is chordal, it turns out that this condition is also sufficient. The following theorem can actually be obtained from Theorem 29 via duality.

Theorem 30. *([56]) Let $\mathcal{G} = (V, E)$ be a graph and let X be a $\mathcal{G}$-partial matrix. Assume that $\mathcal{G}$ is chordal. Then X can be completed into a full $|V| \times |V|$ Hermitian positive semidefinite matrix if, and only if, $X[\mathcal{C}, \mathcal{C}] \succeq 0$ for all maximal cliques $\mathcal{C}$ of $\mathcal{G}$.*

5.4 Main theorem

Let G be a finite abelian group and let $\mathbb{C}^G$ be the space of complex-valued functions on G. Given a nonnegative function $f : G \to \mathbb{R}_+$, a *sum-of-squares certificate* for f

takes the form

$$f(x) = \sum_{k=1}^{K} |f_k(x)|^2 \quad \forall x \in G, \tag{5.16}$$

where $f_1, \ldots, f_K \in \mathbb{C}^G$.

It is well-known in the literature on polynomial optimization (see e.g., [85, 83, 68]), and as we saw in the proof of Theorem 5 (Chapter 2), that the existence of sum-of-squares certificates can be expressed in terms of the existence of a certain positive semidefinite matrix called a *Gram* matrix for f. This connection between sum-of-squares certificates and positive semidefinite matrices will be important, and so we recall this connection more formally in the next proposition.

Proposition 9. *Let $n = |G|$ and let $b_1, \ldots, b_n$ be a basis for $\mathbb{C}^G$. Let $f : G \to \mathbb{R}$ be a real-valued function on G. Then f has a sum-of-squares representation (5.16), if, and only if, there exists a $n \times n$ Hermitian positive semidefinite matrix Q such that*

$$f(x) = [b(x)]^* Q [b(x)] = \sum_{1 \leq i,j \leq n} Q_{ij} \overline{b_i(x)} b_j(x) \quad \forall x \in G \tag{5.17}$$

where $[b(x)] := [b_i(x)]_{i=1,\ldots,n} \in \mathbb{C}^n$. If (5.17) holds where Q is Hermitian positive semidefinite, we say that Q is a Gram matrix *for f in the basis $b_1, \ldots, b_n$.*

Proof. Assume first that f is a sum of squares, i.e., $f(x) = \sum_{k=1}^{K} |f_k(x)|^2$. Since $(b_1, \ldots, b_n)$ forms a basis of $\mathbb{C}^G$ we can write $f_k(x) = \sum_{i=1}^{n} \overline{a_{ki}} b_i(x)$ for some coefficients $a_{ki} \in \mathbb{C}$. Note that $|f_k(x)|^2 = \sum_{1 \leq i,j \leq n} a_{ki} \overline{a_{kj}} \overline{b_i(x)} b_j(x)$ and thus $f(x) = \sum_k |f_k(x)|^2 = \sum_{1 \leq i,j \leq n} Q_{i,j} \overline{b_i(x)} b_j(x)$ where Q is the Hermitian matrix defined by: $Q_{i,j} = \sum_k a_{ki} \overline{a_{kj}}$. Note that Q is positive semidefinite since it has the form $Q = \sum_k a_k a_k^*$ where a_k is the vector $(a_k)_i = a_{ki}$.

We now show the converse. Assume f can be written as (5.17). Since Q is positive semidefinite, we can find vectors a_k such that $Q = \sum_{k=1}^{K} a_k a_k^*$. If we define f_k to be the function $f_k(x) = \sum_{i=1}^{n} \overline{a_{ki}} b_i(x)$ then we can verify that $f = \sum_{k=1}^{K} |f_k|^2$. □

Given $y \in G$ define the Dirac function δ_y at y by

$$\delta_y(x) = \begin{cases} 1 & \text{if } x = y \\ 0 & \text{else.} \end{cases}$$

Then it is easy to see that every nonnegative function on G is a sum of squares.

Proposition 10. *Any nonnegative function f on G has a sum-of-squares certificate as*

$$f(x) = \sum_{y \in G} |\sqrt{f(y)} \delta_y(x)|^2 \quad \forall x \in G. \tag{5.18}$$

Said differently, a nonnegative function f is a sum-of-squares because if we pick $b_1, \ldots, b_n$ to be the basis of Dirac functions, then f satisfies Equation (5.17) where Q is the diagonal matrix consisting of the values taken by f on G.

Since we are working with functions on a finite abelian group G, it is more natural (and more beneficial, as we see later) to look at sum-of-squares representation in the basis of *characters*. One reason for this is that typically the functions f we are interested in have a small support in the basis of characters and in this case one can hope to find a sum-of-squares decomposition which also only involves a small number of characters. The next proposition is simply a change-of-basis in the formula (5.18).

Proposition 11. *Let $f : G \to \mathbb{R}$ and assume that f is nonnegative, i.e., $f(x) \geq 0$ for all $x \in G$. Define the Hermitian matrix $Q \in \mathbf{H}^{\widehat{G}}$ (indexed by characters $\chi \in \widehat{G}$) by*

$$Q_{\chi,\chi'} = \widehat{f}(\overline{\chi}\chi'). \tag{5.19}$$

Then Q is positive semidefinite and we have, for any $x \in G$,

$$f(x) = \frac{1}{|G|}[\chi(x)]^* Q[\chi(x)] = \frac{1}{|G|} \sum_{\chi,\chi' \in \widehat{G}} Q_{\chi,\chi'} \overline{\chi(x)} \chi'(x) \tag{5.20}$$

where $[\chi(x)] := [\chi(x)]_{\chi \in \widehat{G}} \in \mathbb{C}^{\widehat{G}}$.

Proof. Consider the matrix $X = [\chi(x)]_{x \in G, \chi \in \widehat{G}}$ where rows are indexed by elements $x \in G$ and columns are indexed by characters $\chi \in \widehat{G}$. Since the characters form an orthonormal basis of $\mathbb{C}^G$ for the inner product (5.14), this means that the matrix $\frac{1}{\sqrt{|G|}}X$ is a unitary matrix. Note that we can rewrite the definition (5.19) of Q in matrix terms as

$$Q = \frac{1}{|G|} X^* \operatorname{diag}([f(x)]_{x \in G}) X,$$

where $\operatorname{diag}([f(x)]_{x \in G})$ is the diagonal matrix with the values $f(x)$ on the diagonal. This shows that the eigenvalues of Q are the values $\{f(x), x \in G\}$, and thus Q is positive semidefinite. Since $\frac{1}{\sqrt{|G|}}X$ is unitary we also get that

$$\operatorname{diag}([f(x)]_{x \in G}) = \frac{1}{|G|} X Q X^*.$$

which, when evaluated at the diagonal entries, is exactly Equation (5.20). □

Example 14. We now include a simple example to illustrate the previous theorem. Let $G = \mathbb{Z}_6$ and consider the function

$$f(x) = 1 - \frac{1}{2}(\chi_1(x) + \chi_{-1}(x)) = 1 - \cos(2\pi x/6) \quad \forall x \in \mathbb{Z}_6. \tag{5.21}$$

Clearly $f(x) \geq 0$ for all $x \in \mathbb{Z}_6$. Also note that $\widehat{f}(0) = 1$, $\widehat{f}(1) = \widehat{f}(-1) = -1/2$ and $\widehat{f}(k) = 0$ for all $k \notin \{-1, 0, 1\}$. The matrix Q defined in (5.19) associated to this

function f takes the form

$$Q = \begin{bmatrix} 1 & -1/2 & 0 & 0 & 0 & -1/2 \\ -1/2 & 1 & -1/2 & 0 & 0 & 0 \\ 0 & -1/2 & 1 & -1/2 & 0 & 0 \\ 0 & 0 & -1/2 & 1 & -1/2 & 0 \\ 0 & 0 & 0 & -1/2 & 1 & -1/2 \\ -1/2 & 0 & 0 & 0 & -1/2 & 1 \end{bmatrix}. \tag{5.22}$$

$\Diamond$

We are now interested in nonnegative functions $f : G \to \mathbb{R}$ with Fourier support on a subset $\mathcal{S} \subseteq \widehat{G}$, i.e., $\widehat{f}(\chi) = 0$ for all $\chi \notin \mathcal{S}$. For such functions we are interested in finding *Fourier-sparse* sum-of-squares certificates for f, i.e., we are interested in finding a set $\mathcal{T} \subseteq \widehat{G}$ such that any nonnegative function f with Fourier support $\mathcal{S}$ has a sum-of-squares certificate of the form

$$f = \sum_{k=1}^{K} |f_k|^2 \quad \text{where} \quad \operatorname{supp} \widehat{f_k} \subseteq \mathcal{T} \quad \forall k = 1, \dots, K. \tag{5.23}$$

The main idea to obtain such a "sparse" sum-of-squares certificate of f is to exploit the sparsity of the Gram matrix Q from Proposition 11. Indeed, note that if $\operatorname{supp} \widehat{f} = \mathcal{S}$, then the Gram matrix Q of Proposition 11 has sparsity pattern given by

$$Q_{\chi,\chi'} \neq 0 \Leftrightarrow \overline{\chi}\chi' \in \mathcal{S}.$$

In other words, the sparsity structure of Q is given by the *Cayley graph* $\operatorname{Cay}(\widehat{G}, \mathcal{S})$. Recall the definition of a Cayley graph.

Definition 15. Let H be a group and let $\mathcal{S} \subset H$ be a subset of H that is symmetric, i.e., $x \in \mathcal{S} \Rightarrow x^{-1} \in \mathcal{S}$. The Cayley graph $\operatorname{Cay}(H, \mathcal{S})$ is the graph where vertices are the elements of the group H, and where two distinct vertices $x, y \in H$ are connected by an edge if $x^{-1}y \in \mathcal{S}$ (or $y^{-1}x \in \mathcal{S}$, which is the same since $\mathcal{S}$ is symmetric).

Remark 16. We do not require the set $\mathcal{S}$ to be a generator for the group H and hence the graph $\operatorname{Cay}(H, \mathcal{S})$ may be disconnected. Also observe that the set $\mathcal{S} = \operatorname{supp} \widehat{f}$ in our case is symmetric since f is real-valued; indeed when f is real-valued we have $\widehat{f}(\overline{\chi}) = \overline{\widehat{f}(\chi)}$ for all $\chi \in \widehat{G}$ and thus $\chi \in \operatorname{supp} \widehat{f} \Rightarrow \overline{\chi} \in \operatorname{supp} \widehat{f}$.

To obtain a set $\mathcal{T} \subseteq \widehat{G}$ such that (5.23) holds for all functions f with Fourier support $\mathcal{S}$ we will study *chordal covers* of the graph $\operatorname{Cay}(\widehat{G}, \mathcal{S})$. We now introduce the key definition of *Fourier support* for a graph with vertices $\widehat{G}$.

Definition 16. Let Γ be a graph with vertices $\widehat{G}$. We say that Γ has *Fourier support* $\mathcal{T} \subseteq \widehat{G}$ if for any maximal clique $\mathcal{C}$ of Γ there exists $\chi_{\mathcal{C}} \in \widehat{G}$ such that $\chi_{\mathcal{C}}\mathcal{C} \subseteq \mathcal{T}$ (where $\chi_{\mathcal{C}}\mathcal{C} := \{\chi_{\mathcal{C}}\chi : \chi \in \mathcal{C}\}$ is the translation of $\mathcal{C}$ by $\chi_{\mathcal{C}}$).

Note that one can also state the definition of Fourier support of Γ in terms of *equivalence classes* of cliques. Given a subset $\mathcal{C} \subseteq \widehat{G}$ define the *equivalence class* of

$\mathcal{C}$ to be all the subsets of $\widehat{G}$ that can be obtained from $\mathcal{C}$ by translation, i.e., it is the set $[\mathcal{C}] := \{\chi\mathcal{C} : \chi \in \widehat{G}\}$. Using this terminology, the graph Γ has Fourier support $\mathcal{T}$ if for any maximal clique $\mathcal{C}$ of Γ there is at least one representative from $[\mathcal{C}]$ that is contained in $\mathcal{T}$.

We are now ready to state and prove our main theorem (the theorem below was stated as Theorem 25 in the introduction and we reuse the same numbering here since it is just a restatement).

Theorem 25. *Let $\mathcal{S}$ be a symmetric subset of $\widehat{G}$ and assume that $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ has a chordal cover Γ with Fourier support $\mathcal{T} \subseteq \widehat{G}$. Then any nonnegative function with Fourier support $\mathcal{S}$ admits a sum-of-squares certificate with Fourier support $\mathcal{T}$.*

Proof. Let $f : G \to \mathbb{R}$ be a nonnegative function with Fourier support $\mathcal{S}$. Let Q be the Gram matrix (5.19) associated to the sum-of-squares representation of f in the basis of characters. We saw that Q is sparse according to the Cayley graph $\mathrm{Cay}(\widehat{G}, \mathcal{S})$. Since Γ is a cover of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$, Q is also sparse according to Γ. Thus, since Γ is chordal, using Theorem 29 we can find a decomposition of Q as

$$Q = \sum_{\mathcal{C}} Q_{\mathcal{C}} \tag{5.24}$$

where the sum is over the maximal cliques $\mathcal{C}$ of Γ and where each $Q_{\mathcal{C}}$ is a positive semidefinite matrix supported on $\mathcal{C}$. Note that Equation (5.24) implies that for all $x \in G$,

$$[\chi(x)]^* Q [\chi(x)] = \sum_{\mathcal{C}} [\chi(x)]^* Q_{\mathcal{C}} [\chi(x)].$$

Since $f(x) = [\chi(x)]^* Q [\chi(x)]/|G|$ the above equation says that

$$f(x) = \sum_{\mathcal{C}} f_{\mathcal{C}}(x)$$

where we let $f_{\mathcal{C}}(x) := [\chi(x)]^* Q_{\mathcal{C}} [\chi(x)]/|G|$. Since $Q_{\mathcal{C}}$ is positive semidefinite and supported on $\mathcal{C}$, this means that each $f_{\mathcal{C}}(x)$ is a sum-of-squares of functions with Fourier support $\mathcal{C} \subseteq \widehat{G}$, i.e.,

$$f_{\mathcal{C}} = \sum_k |f_{\mathcal{C},k}|^2$$

where $\operatorname{supp} \widehat{f_{\mathcal{C},k}} \subseteq \mathcal{C}$.

According to Definition 16, we know that there exist $\chi_{\mathcal{C}} \in \widehat{G}$ for each maximal clique $\mathcal{C}$ of Γ such that $\chi_{\mathcal{C}}\mathcal{C} \subseteq \mathcal{T}$. Now, observe that

$$f = \sum_{\mathcal{C}} f_{\mathcal{C}} = \sum_{\mathcal{C}} \sum_k |f_{\mathcal{C},k}|^2 \overset{(i)}{=} \sum_{\mathcal{C}} \sum_k |\chi_{\mathcal{C}} f_{\mathcal{C},k}|^2 \overset{(ii)}{=} \sum_{\mathcal{C}} \sum_k |\widetilde{f}_{\mathcal{C},k}|^2$$

where in (i) we used the fact that $|\chi_{\mathcal{C}}|^2 = 1$ and in (ii) we let $\widetilde{f}_{\mathcal{C},k} = \chi_{\mathcal{C}} f_{\mathcal{C},k}$ which has Fourier support $\chi_{\mathcal{C}}\mathcal{C} \subseteq \mathcal{T}$. Thus we have shown that f is a sum-of-squares of functions with Fourier support $\mathcal{T}$. □

Example 15. Let $G = \mathbb{Z}_6$ and let $\mathcal{S} = \{-1, 0, 1\} \subset \widehat{\mathbb{Z}_6}$. We will use the previous theorem to show that any nonnegative function on $\mathbb{Z}_6$ with Fourier support $\mathcal{S} = \{-1, 0, 1\}$ is a sum-of-squares of functions with Fourier support $\mathcal{T} = \{-1, 0, 1, 3\} \subseteq \widehat{\mathbb{Z}_6}$. The Cayley graph $\mathrm{Cay}(\widehat{\mathbb{Z}_6}, \{-1, 0, 1\})$ is the cycle graph on 6 nodes shown in Figure 5-2(left). Clearly the graph is not chordal since the cycle $0, 1, \ldots, 5$ has no chord. Figure 5-2(right) shows a chordal cover Γ of $\mathrm{Cay}(\widehat{\mathbb{Z}_6}, \{-1, 0, 1\})$ where the maximal cliques are

$$\mathcal{C}_1 = \{0, 1, 3\}, \;\; \mathcal{C}_2 = \{1, 2, 3\}, \;\; \mathcal{C}_3 = \{3, 4, 5\}, \;\; \mathcal{C}_4 = \{0, 3, 5\}.$$

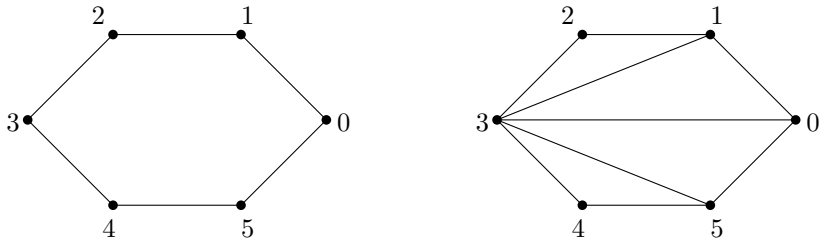

$\mathrm{Cay}(\widehat{\mathbb{Z}_6}, \{-1, 0, 1\})$ Γ

Figure 5-2: Left: The Cayley graph $\mathrm{Cay}(\widehat{\mathbb{Z}_6}, \{-1, 0, 1\})$ is the cycle graph on 6 nodes. Right: A chordal cover of the cycle graph, Γ.

Observe that if we translate the clique $\mathcal{C}_2 = \{1, 2, 3\}$ by -2 we get $\{-1, 0, 1\}$ and similarly if we translate the clique $\{3, 4, 5\}$ by -4 we also get $\{-1, 0, 1\}$. Thus by choosing

$$\chi_{\mathcal{C}_1} = 0, \;\; \chi_{\mathcal{C}_2} = -2, \;\; \chi_{\mathcal{C}_3} = -4, \;\; \chi_{\mathcal{C}_4} = 0$$

we get that $\chi_{\mathcal{C}} + \mathcal{C} \subseteq \{-1, 0, 1, 3\}$ for all maximal cliques $\mathcal{C}$ of Γ (we used the fact that $5 = -1$ in $\mathbb{Z}_6$). In other words we have shown that Γ is a chordal cover of $\mathrm{Cay}(\widehat{\mathbb{Z}_6}, \{-1, 0, 1\})$ with Fourier support $\{-1, 0, 1, 3\}$. Thus by Theorem 25, this means that any nonnegative function on $\mathbb{Z}_6$ with Fourier support $\{-1, 0, 1\}$ can be written as a sum-of-squares of functions with Fourier support $\{-1, 0, 1, 3\}$. $\lozenge$

5.5 Application 1: cut polytope and Laurent's conjecture

In this section we apply the results of Section 5.4 to the case of nonnegative quadratic forms on the vertices of the hypercube in n dimensions. Dually, the moment polytope of interest in this section is the nth *cut polytope* CUT_n. Our main aim is to establish Laurent's conjecture [72, Conjecture 4] that any nonnegative quadratic form on the vertices of the hypercube in n dimension is a sum of squares of polynomials of degree at most $\lceil n/2 \rceil$.

5.5.1 Quadratic forms on $\{-1,1\}^n$ and the cut polytope

Let $G = \{-1,1\}^n$ be the vertices of the hypercube in dimension n. View G as a group (isomorphic to $\mathbb{Z}_2^n$) under componentwise multiplication. Recall that the characters of G are indexed by subsets $S \in 2^{[n]}$ and are the square-free monomials

$$\chi_S(x) = \prod_{i \in S} x_i \quad \text{for all } x \in G.$$

We focus on characterizing nonnegative quadratic functions on G. These are of particular interest because the problem of maximizing a quadratic form over G i.e.

$$\max_{x \in G} \sum_{1 \le i < j \le n} A_{ij} x_i x_j \tag{5.25}$$

includes, for example, the MAX-CUT problem, which arises when the symmetric matrix A_{ij} is the Laplacian of a (weighted) graph on n vertices. We can solve (5.25) by finding the smallest upper bound on the objective:

$$\min_{\gamma} \quad \gamma \quad \text{s.t.} \quad \gamma - \sum_{1 \le i < j \le n} A_{ij} x_i x_j \ge 0 \quad \text{for all } x \in G. \tag{5.26}$$

If we have a characterization of nonnegative functions on G with Fourier support $\mathcal{S} = \{S \in 2^{[n]} : |S| = 0 \text{ of } |S| = 2\}$ as sums of squares of functions with Fourier support $\mathcal{T} \subseteq \widehat{G}$ then we can solve (5.26) by solving a semidefinite optimization problem of size $|\mathcal{T}|$.

The dual picture to (5.26) is to consider optimization over the moment polytope $\mathcal{M}(\{-1,1\}^n, \mathcal{S} \setminus \{\emptyset\})$, known as the *cut polytope*

$$\mathrm{CUT}_n := \mathcal{M}(\{-1,1\}^n, \mathcal{S} \setminus \{\emptyset\}) = \mathrm{conv}\,\{(x_i x_j)_{1 \le i < j \le n} : (x_1, x_2, \ldots, x_n) \in \{-1,1\}^n\}.$$

We can solve the binary quadratic optimization problem (5.25) by optimizing the linear function defined by A over CUT_n, i.e. by solving the linear program

$$\max_{(\ell_{ij})_{1 \le i < j \le n}} \sum_{i<j} \ell_{ij} A_{ij} \quad \text{s.t.} \quad (\ell_{ij})_{1 \le i < j \le n} \in \mathrm{CUT}_n.$$

If we have a SDP lift of the cut polytope CUT_n of size $|\mathcal{T}|$ then we can solve this optimization problem by solving a semidefinite optimization problem of size $|\mathcal{T}|$.

5.5.2 The associated Cayley graph

To apply the results of Section 5.4 we need to understand the graph $\mathrm{Cay}(\widehat{G}, \mathcal{S})$. In the case $n = 4$ this graph is shown in Figure 5-3. Throughout this section we identify the character $\chi_S \in \widehat{G}$ with the subset $S \subseteq [n]$ that indexes it and work exclusively in the language of subsets. As such, the vertex set of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ is $2^{[n]}$, the collection of all subsets $S \subseteq [n]$. There is an edge between two subsets S, T if and only if $|S \triangle T| = 2$.

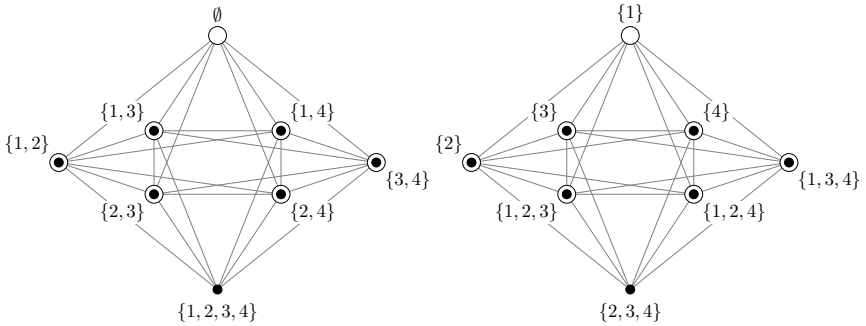

Figure 5-3: The Cayley graph $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ for $G = \{-1, 1\}^4$ and $\mathcal{S} = \{S : |S| = 0 \text{ or } |S| = 2\}$. The two connected components are $\mathcal{T}_{\text{even}}$ (left) and $\mathcal{T}_{\text{odd}}$ (right). The vertices of $\mathcal{T}_{\text{odd}}$ are arranged to correspond to their images in $\mathcal{T}_{\text{even}}$ under the graph automorphism $\phi(S) = \{1\} \triangle S$. We can obtain a chordal cover of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ by forming maximal cliques on the vertices of $\mathcal{T}_{\text{even}}$ marked with filled circles, the vertices of $\mathcal{T}_{\text{even}}$ marked with open circles, and the images in $\mathcal{T}_{\text{odd}}$ of these two cliques under the map ϕ.

This graph is often called the *half-cube graph*.

The group operation on characters is multiplication of functions, which corresponds to taking the symmetric difference of the subsets that index the characters. In other words, if $S, T \subseteq [n]$ then

$$\chi_S(x)\chi_T(x) = \chi_{S \triangle T}(x)$$

where $S \triangle T = (S \setminus T) \cup (T \setminus S)$. As such, there is an action of $\widehat{G}$ on the vertices of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ by $S \cdot T = S \triangle T$. Furthermore if $\mathcal{T} \subseteq 2^{[n]}$ is a subset of the vertices of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ we write $S \triangle \mathcal{T} := \{S \triangle T : T \in \mathcal{T}\}$.

We now record some simple observations that follow directly from the adjacency relation in $\mathrm{Cay}(\widehat{G}, \mathcal{S})$. For convenience of notation, for $k = 0, 1, \ldots, n$ let

$$\mathcal{T}_k = \{S \subseteq [n] : |S| = k\}.$$

Any edge of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ either has both endpoints in $\mathcal{T}_k$ for some k or one endpoint in $\mathcal{T}_k$ and the other in $\mathcal{T}_{k+2}$ for some k. Consequently, $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ has two connected components

$$\mathcal{T}_{\text{even}} = \mathcal{T}_0 \cup \mathcal{T}_2 \cup \cdots \cup \mathcal{T}_{2\lfloor n/2 \rfloor} \quad \text{and} \quad \mathcal{T}_{\text{odd}} = \mathcal{T}_1 \cup \mathcal{T}_3 \cup \cdots \cup \mathcal{T}_{2\lceil n/2 \rceil - 1}.$$

Define a map $\phi : 2^{[n]} \to 2^{[n]}$ by $\phi(S) = \{1\} \triangle S$. Since $|\phi(S) \triangle \phi(T)| = |S \triangle T|$ for all $S, T \in 2^{[n]}$ it follows that ϕ extends to an automorphism of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ that exchanges $\mathcal{T}_{\text{even}}$ and $\mathcal{T}_{\text{odd}}$.

5.5.3 Applying Theorem 25

To apply Theorem 25 from Section 5.4 we need to find a subset $\mathcal{T} \subseteq 2^{[n]}$ of vertices such that $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ has a chordal cover with Fourier support $\mathcal{T}$. The following result explicitly describes such a collection of vertices.

Proposition 12. *The graph* $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ *has a chordal cover with Fourier support*

$$\mathcal{T} = \begin{cases} \mathcal{T}_0 \cup \mathcal{T}_2 \cup \cdots \cup \mathcal{T}_{\lceil n/2 \rceil} & \text{if } \lceil n/2 \rceil \text{ even} \\ \mathcal{T}_1 \cup \mathcal{T}_3 \cup \cdots \cup \mathcal{T}_{\lceil n/2 \rceil} & \text{if } \lceil n/2 \rceil \text{ odd.} \end{cases} \tag{5.27}$$

Proof. We give a detailed proof in Section 5.8.2. □

Example 16. To give the flavor of the proof, we discuss the case $n = 4$. In this case $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ is shown in Figure 5-3. Define Γ to be the graph with vertex set $2^{[4]}$ and with edges between $S, T \in \mathcal{T}_{\text{even}}$ if and only if $||S| - |T|| \leq 2$, and edges between $S, T \in \mathcal{T}_{\text{odd}}$ if and only if $||\phi(S)| - |\phi(T)|| \leq 2$. The graph Γ is chordal, with maximal cliques given by $\mathcal{C}_0 = \mathcal{T}_0 \cup \mathcal{T}_2$, $\mathcal{C}_2 = \mathcal{T}_2 \cup \mathcal{T}_4$, $\phi(\mathcal{C}_0)$, and $\phi(\mathcal{C}_2)$. The vertices in cliques $\mathcal{C}_0$ and $\mathcal{C}_2$ are indicated by open and filled circles respectively in Figure 5-3. (The vertices in cliques $\phi(\mathcal{C}_0)$ and $\phi(\mathcal{C}_2)$ are similarly marked.) If $\mathcal{T} = \mathcal{T}_0 \cup \mathcal{T}_2$ then we can see that Γ is a chordal cover of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ with Fourier support $\mathcal{T}$ by observing that $\emptyset \triangle \mathcal{C}_0 \subseteq \mathcal{T}$, $\{1,2,3,4\} \triangle \mathcal{C}_2 \subseteq \mathcal{T}$, $\phi(\emptyset) \triangle \phi(\mathcal{C}_0) \subseteq \mathcal{T}$ and $\phi(\{1,2,3,4\}) \triangle \phi(\mathcal{C}_2) \subseteq \mathcal{T}$. ◊

Laurent's conjecture follows directly from Proposition 12 and Theorem 25.

Theorem 26. *Suppose* $f(x) = A_\emptyset + \sum_{1 \leq i < j \leq n} A_{ij} x_i x_j$ *is nonnegative on* $G = \{-1, 1\}^n$. *Then there is a collection* $(h_k)_{k=1}^{|\mathcal{T}|}$ *of functions* $h_k : G \to \mathbb{R}$ *each with Fourier support* $\mathcal{T}$ *(defined in (5.27)) such that*

$$f(x) = \sum_{k=1}^{|\mathcal{T}|} h_k(x)^2.$$

Consequently, any nonnegative quadratic form on G *is a sum of squares of functions of degree at most* $\lceil n/2 \rceil$.

Proof. The first assertion follows directly from Proposition 12 and Theorem 25. The second assertion holds simply because every function with Fourier support $\mathcal{T}$ has degree at most $\lceil n/2 \rceil$. □

The dual version of this result gives a SDP lift of the cut polytope of size $|\mathcal{T}|$. It follows directly from Proposition 12 and Theorem 25D, and the observation that in this case all the characters are real-valued.

Corollary 2. *The cut polytope* CUT_n *has a real SDP lift of size* $|\mathcal{T}|$ *given by*

$$CUT_n = \Big\{ \ell \in \mathbb{R}^{\mathcal{S} \setminus \emptyset} : \exists y \in \mathbb{R}^{\mathcal{T} \triangle \mathcal{T}} \;\; s.t. \;\; y_\emptyset = 1, \;\; y_{\{i,j\}} = \ell_{\{i,j\}} \text{ for } 1 \leq i < j \leq n \\ \text{and } [y_{S \triangle T}]_{S,T \in \mathcal{T}} \succeq 0 \Big\}$$

where $\mathcal{T}$ is defined in (5.27).

In the language used in [72], Corollary 2 simply expresses that $Q_{\lceil n/2 \rceil} = \mathrm{CUT}_n$ where Q_k is the k'th Lasserre semidefinite relaxation of CUT_n.

5.6 Application 2: trigonometric cyclic polytopes

In this section we apply the results of Section 5.4 to the case where $G = \mathbb{Z}_N$ is the (additive) group of integers modulo N. As we will see, this will allow us to obtain a positive semidefinite lift of size $O(d \log(N/d))$ for the regular trigonometric cyclic polytope with N vertices of degree d, when d divides N.

Recall from Section 5.3, that the characters of $\mathbb{Z}_N$ are indexed by $k \in \mathbb{Z}_N$ and are given by

$$\chi_k(x) = e^{2i\pi kx/N} \quad \forall x \in \mathbb{Z}_N.$$

Thus the Fourier decomposition of a function $f : \mathbb{Z}_N \to \mathbb{C}$ is given by

$$f(x) = \sum_{k \in \mathbb{Z}_N} \widehat{f}(k) e^{2i\pi kx/N}.$$

Furthermore, we say that a function f has degree d if it has Fourier support $\{-d, -(d-1), \ldots, d-1, d\}$.

5.6.1 The case $\mathcal{S} = \{-1, 0, 1\}$: the cycle graph

In this section we are interested in obtaining Fourier-sparse sum-of-squares certificates for functions of degree 1 on $\mathbb{Z}_N$, i.e., functions with Fourier support $\mathcal{S} = \{-1, 0, 1\}$. Note that the moment polytope in this case is nothing but the regular N-gon (up to a linear transformation). Indeed $\mathcal{M}(\mathbb{Z}_N, \{-1, 1\})$ is by definition

$$\mathcal{M}(\mathbb{Z}_N, \{-1, 1\}) = \mathrm{conv}\left\{(e^{2i\pi x/N}, e^{-2i\pi x/N}) : x \in \mathbb{Z}_N\right\}.$$

Under the invertible $\mathbb{R}$-linear map $(z, \overline{z}) \mapsto (\mathrm{Re}(z), \mathrm{Im}(z))$, we see that $\mathcal{M}(\mathbb{Z}_N, \{-1, 0, 1\})$ is linearly isomorphic to the regular N-gon in $\mathbb{R}^2$.

To obtain Fourier-sparse sum-of-squares certificates for nonnegative functions of degree 1 we are going to study the Cayley graph $\mathrm{Cay}(\widehat{\mathbb{Z}_N}, \{-1, 0, 1\})$. Note that this is simply the cycle graph on N vertices, which we will denote by C_N for simplicity. The object of this section is to show that this graph admits a chordal cover with small Fourier support.

Theorem 31. *Let N be a positive integer greater than 2. Then the cycle graph C_N has a chordal cover with Fourier support $\mathcal{T} \subseteq \widehat{\mathbb{Z}_N}$ where $|\mathcal{T}| \leq 3 \log_2 N$. More precisely the set $\mathcal{T}$ can be described explicitly as follows: Let $k_1 < k_2 < \cdots < k_l$ be the positions of the nonzero digits in the binary expansion of N so that $N = \sum_{j=1}^{l} 2^{k_j}$. Let k be the largest integer such that $2^k < N$ (i.e., $k = k_l - 1$ if N is a power of two and $k = k_l$*

otherwise). Then the set $\mathcal{T}$ is given by

$$\mathcal{T} = \{0\} \cup \{-2^i, 2^i, i = 0, \ldots, k-1\} \cup \left\{ \sum_{j=1}^{i} 2^{k_j}, i = 1, \ldots, l-2 \right\}. \tag{5.28}$$

Proof. The chordal cover is constructed by induction on N, see Section 5.8.1 for the details. Figure 5-4 shows the chordal cover for $N = 8$ and $N = 16$. □

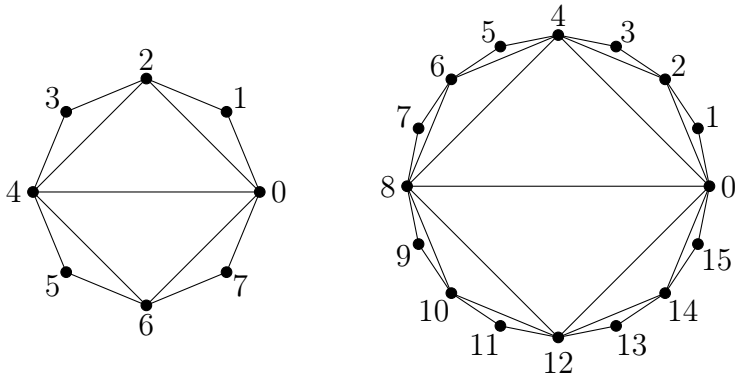

Figure 5-4: Chordal cover of the 8-cycle with Fourier support $\mathcal{T} = \{-2, -1, 0, 1, 2\}$ and of the 16-cycle with Fourier support $\mathcal{T} = \{-4, -2, -1, 0, 1, 2, 4\}$.

If we combine the previous theorem with Theorem 25, we get that any nonnegative degree-1 function on $\mathbb{Z}_N$ has a sum-of-squares certificate with Fourier support $\mathcal{T}$ where $|\mathcal{T}| \leq 3 \log N$. Note that this corresponds to Theorem 27 from the introduction for the case $d = 1$. Dually, this allows us to obtain a Hermitian positive semidefinite lift of the regular N-gon of size $|\mathcal{T}| \leq 3 \log N$.

In Section 5.1 we showed that the $N = 2^n$-gon admits a positive semidefinite lift of size $2n - 1$. In fact we showed in Proposition 8 that any linear function on $\mathbb{Z}_N$ that is nonnegative can be written as a sum-of-squares of functions with Fourier support $\{0\} \cup \{\pm 2^i, i = 0, \ldots, n-2\}$. Note that this is the same Fourier support that we get if we plug $N = 2^n$ in (5.28). Thus Theorem 31 generalizes the result of Section 5.1 to arbitrary N.

5.6.2 Degree d functions: powers of cycle graph

In this section we are interested in functions of degree d on $\mathbb{Z}_N$ where d divides N. We show how to construct a chordal cover of the associated Cayley graph $\mathrm{Cay}(\widehat{\mathbb{Z}_N}, \mathcal{S})$ using the chordal cover of the cycle graph constructed in the previous section. This allows us to show that any nonnegative function on $\mathbb{Z}_N$ of degree d has a sum-of-squares certificate with Fourier support of size most $3d \log(N/d)$.

Constructing a chordal cover of the Cayley graph

We start by recalling the definition of the power of a graph.

Definition 17. Let $\mathcal{G} = (V, E)$ be a graph. Given $d \in \mathbb{N}$, the d'th power of $\mathcal{G}$ is the graph $\mathcal{G}^d = (V, E^d)$ where two vertices $u, v \in V$ are connected by an edge if there is a path of length $\leq d$ connecting u and v in $\mathcal{G}$.

It is not difficult to see that the Cayley graph $\mathrm{Cay}(\widehat{\mathbb{Z}_N}, \mathcal{S})$ when $\mathcal{S} = \{-d, \ldots, d\}$ is the d'th *power* of the cycle graph $\mathrm{Cay}(\widehat{\mathbb{Z}_N}, \{-1, 1\})$. Following this observation, we will use the symbol C_N^d to denote the Cayley graph $\mathrm{Cay}(\widehat{\mathbb{Z}_N}, \{-d, \ldots, d\})$. Figure 5-5(left) shows the graph C_N^d for $N = 8$ and $d = 2$.

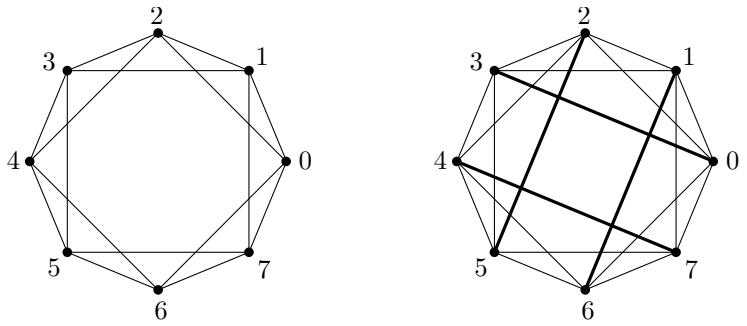

C_8^2 $\qquad$ $C_4 \boxtimes K_2$

Figure 5-5: Left: The second power of the cycle graph on 8 nodes: two nodes are connected by an edge if their distance in the cycle graph is at most 2. Right: The graph $C_4 \boxtimes K_2$. Note that $C_8^2 \subset C_4 \boxtimes K_2$. The edges in $C_4 \boxtimes K_2$ that are not in C_8^2 are indicated with a heavier line.

To construct a chordal cover of C_N^d we will use the chordal cover of the cycle graph C_N constructed in the previous section. For this, we need the following definition of *strong product* of graphs.

Definition 18. Given graphs $\mathcal{G} = (V, E)$ and $\mathcal{G}' = (V', E')$ define the *strong product* of $\mathcal{G}$ and $\mathcal{G}'$, denoted $\mathcal{G} \boxtimes \mathcal{G}'$ to be the graph with vertex set $V \times V'$ and where two vertices $(u, u') \in V \times V'$ and $(v, v') \in V \times V'$ are connected if

$$\begin{aligned} &(u = v \text{ and } \{u', v'\} \in E') \\ \text{or } &(\{u, v\} \in E \text{ and } u' = v') \\ \text{or } &(\{u, v\} \in E \text{ and } \{u', v'\} \in E'). \end{aligned}$$

Remark 17. An important special case is when one of the graphs, say $\mathcal{G}'$, is a complete graph $\mathcal{G}' = K_m$. In this case two distinct vertices (u, u') and (v, v') in $\mathcal{G} \boxtimes K_m$ are connected if either $u = v$ or $\{u, v\} \in E(\mathcal{G})$.

Given two graphs $\mathcal{G} = (V, E)$ and $\mathcal{G}' = (V, E')$ with the same vertex set V we say that $\mathcal{G}'$ covers $\mathcal{G}$ and we write $\mathcal{G} \subseteq \mathcal{G}'$ if $E \subseteq E'$. We use the following observation to construct a chordal cover of C_N^d.

Proposition 13. *Let N and d be two integers and assume that d divides N. Let C_N^d be the d'th power of the cycle graph C_N and let $C_{N/d}$ be the cycle graph on N/d nodes. Then*

$$C_N^d \subseteq C_{N/d} \boxtimes K_d. \tag{5.29}$$

Proof. To show the inclusion (5.29) we first need to identify the vertices of C_N^d with those of $C_{N/d} \boxtimes K_d$. Note that the vertex set of C_N^d can be identified with $\mathbb{Z}_N$ and the vertex set of $C_{N/d}$ can be identified with $\mathbb{Z}_{N/d}$. We also identify the vertices of K_d with $\{0, \dots, d-1\}$. By definition of $\boxtimes$, the vertices of $C_{N/d} \boxtimes K_d$ are $\mathbb{Z}_{N/d} \times \{0, \dots, d-1\}$. Consider the map

$$\phi : \mathbb{Z}_{N/d} \times \{0, \dots, d-1\} \to \mathbb{Z}_N, \quad \phi(q, r) = qd + r. \tag{5.30}$$

This map is well-defined and gives a bijection between $\mathbb{Z}_{N/d} \times \{0, \dots, d-1\}$ and $\mathbb{Z}_N$. The map ϕ thus identifies vertices of C_N^d with those of $C_{N/d} \boxtimes K_d$.

We now show that, with this identification, inclusion (5.29) holds. We need to show that if $i, i' \in \mathbb{Z}_N$ are connected in C_N^d (i.e., $i - i' \in \{-d, \dots, d\}$) then necessarily (q, r) and (q', r') are connected in $C_{N/d} \boxtimes K_d$ (i.e., $q - q' \in \{-1, 0, 1\}$), where (q, r) and (q', r') are such that $i = \phi(q, r)$ and $i' = \phi(q', r')$. Consider for $q \in \mathbb{Z}_{N/d}$ the set of vertices of C_N^d given by $V_q = \{\phi(q, r) : r = 0, \dots, d-1\} \subset \mathbb{Z}_N$. Note that $(V_q)_{q \in \mathbb{Z}_{N/d}}$ forms a partition of the vertex set of C_N^d and that $|V_q| = d$ for all q (for example if $N = 8$ and $d = 2$ (Figure 5-5) $V_0 = \{0, 1\}, V_1 = \{2, 3\}, V_2 = \{4, 5\}, V_4 = \{6, 7\}$). It is easy to see that if i and i' are two adjacent vertices of C_N^d, then i and i' must be in the same group (i.e., $i, i' \in V_q$) or in adjacent group (i.e., $i \in V_q$ and $i' \in V_{q+1}$ or vice-versa). In other words this means that $q - q' \in \{-1, 0, 1\}$ which means that (q, r) and (q', r') are connected in $C_{N/d} \boxtimes K_d$. □

The previous proposition gives a natural way to construct a chordal cover of C_N^d from that of $C_{N/d}$. Indeed if Γ is a chordal cover of $C_{N/d}$ then one can show that $\Gamma \boxtimes K_d$ is a chordal cover of C_N^d and one can also characterize the maximal cliques of $\Gamma \boxtimes K_d$ in terms of those of Γ. This is the object of the next proposition.

Proposition 14. *Let $\mathcal{G} = (V, E)$ be a graph and d be any integer.*

1. *If $\mathcal{G}'$ is such that $\mathcal{G} \subseteq \mathcal{G}'$ then $\mathcal{G} \boxtimes K_d \subseteq \mathcal{G}' \boxtimes K_d$.*
2. *If $\mathcal{G}$ is chordal then $\mathcal{G} \boxtimes K_d$ is chordal.*
3. *All the maximal cliques of $\mathcal{G} \boxtimes K_d$ have the form $\mathcal{C} \times K_d$ where $\mathcal{C}$ is a maximal clique of $\mathcal{G}$.*

Proof.

1. The first point is clear from the definition of $\boxtimes$.
2. Let $(u_1, v_1) \dots (u_l, v_l)$ be a cycle in $\mathcal{G} \boxtimes K_d$ of length $l \geq 4$ where $(u_l, v_l) = (u_1, v_1)$. If there exists $i \in \{1, \dots, l-1\}$ such that $u_i = u_{i+1}$ then the edge $\{(u_i, v_i), (u_{i+2}, v_{i+2})\}$ is a chord of the cycle. Otherwise note that $u_1 \dots u_l$ is a cycle in $\mathcal{G}$ of length ≥ 4. Since $\mathcal{G}$ is chordal there is $1 \leq i, j \leq l-1$ with $j - i \geq 2$ such that $\{u_i, u_j\} \in E$. In this case the edge $\{(u_i, v_i), (u_j, v_j)\}$ is a chord of the cycle.
3. The third property easily follows from the fact that if $\mathcal{C} = \{(u_i, v_i), i = 1, \dots, k\}$ is a clique in $\mathcal{G} \boxtimes K_d$ then $\{u_i, i = 1, \dots, k\} \subseteq V$ is a clique in $\mathcal{G}$.

□

We can now use the chordal cover of the cycle graph constructed in the previous section to obtain a chordal cover of C_N^d.

Proposition 15. *Let N and d be two integers and assume that d divides N. If $C_{N/d}$ has a chordal cover with Fourier support $\mathcal{T} \subseteq \mathbb{Z}_{N/d}$, then C_N^d has a chordal cover with Fourier support*

$$\mathcal{T}' = \{dk + r : k \in \mathcal{T}, r \in \{0, \dots, d-1\}\} \tag{5.31}$$

and $|\mathcal{T}'| \leq d \cdot |\mathcal{T}|$.

Proof. Let Γ be a chordal cover $C_{N/d}$ with Fourier support $\mathcal{T}$. By definition, this means that for any maximal clique $\mathcal{C}$ of $C_{N/d}$, there is $k_{\mathcal{C}} \in \mathbb{Z}_{N/d}$ such that $k_{\mathcal{C}} + \mathcal{C} \subseteq \mathcal{T}$.

By Proposition 14, we know that $\Gamma \boxtimes K_d$ is a chordal cover of C_N^d. Let $\mathcal{C}'$ be a maximal clique of $\Gamma \boxtimes K_d$. By Proposition 14, we know that there exists $\mathcal{C}$ maximal clique of Γ such that $\mathcal{C}' = \mathcal{C} \times K_d = \{dq + r : q \in \mathcal{C}, r \in \{0, \dots, d-1\}\}$. Define $k_{\mathcal{C}'} = dk_{\mathcal{C}} \in \mathbb{Z}_N$ and note that

$$\begin{aligned} k_{\mathcal{C}'} + \mathcal{C}' &= \{dk_{\mathcal{C}} + dq + r : q \in \mathcal{C}, r \in \{0, \dots, d-1\}\} \\ &= \{(k_{\mathcal{C}} + q)d + r : q \in \mathcal{C}, r \in \{0, \dots, d-1\}\} \subseteq \mathcal{T}', \end{aligned}$$

where the last inclusion follows from the fact that $k_{\mathcal{C}} + q \in \mathcal{T}$ whenever $q \in \mathcal{C}$. We have thus shown that for any maximal clique $\mathcal{C}'$ of $\Gamma \boxtimes K_d$, there is $k_{\mathcal{C}'} \in \mathbb{Z}_N$ such that $k_{\mathcal{C}'} + \mathcal{C}' \subseteq \mathcal{T}'$. Thus this shows that $\Gamma \boxtimes K_d$ is a chordal cover of C_N^d with Fourier support $\mathcal{T}'$. □

Combining Proposition 15 and the chordal cover of the cycle graph from Theorem 31 we get the following corollary.

Corollary 3. *Let N and d be two integers and assume that d divides N. Then the graph C_N^d has a chordal cover with Fourier support $\mathcal{T} \subset \widehat{\mathbb{Z}_N}$ where $|\mathcal{T}| \leq 3d \log(N/d)$.*

Using Theorem 25, this proves Theorem 27 from the introduction concerning nonnegative functions on $\mathbb{Z}_N$ of degree d.

Theorem 27. *Let N and d be two integers and assume that d divides N. Then there exists $\mathcal{T} \subseteq \mathbb{Z}_N$ with $|\mathcal{T}| \leq 3d \log(N/d)$ such that any nonnegative function on $\mathbb{Z}_N$ of degree at most d has a sum-of-squares certificate with Fourier support $\mathcal{T}$.*

Figure 5-6 shows the chordal cover of C_{16}^2 obtained by applying Theorem 31 to C_8 and applying the strong graph product with K_2.

Cyclic polytopes

Observe that the moment polytope for $G = \mathbb{Z}_N$ and $\mathcal{S} = \{-d, \dots, d\}$ is

$$\mathcal{M}(\mathbb{Z}_N, \{-d, \dots, d\}) = \operatorname{conv}\left\{ (e^{2i\pi kx/N})_{k=-d,\dots,d} : x \in \mathbb{Z}_N \right\} \subset \mathbb{C}^{2d}.$$

This polytope is affinely isomorphic to the regular trigonometric cyclic polytope

$$TC(N, 2d) = \operatorname{conv}\left\{ TM(2\pi x/N) : x = 0, 1, \dots, N-1 \right\} \subset \mathbb{R}^{2d}, \tag{5.32}$$

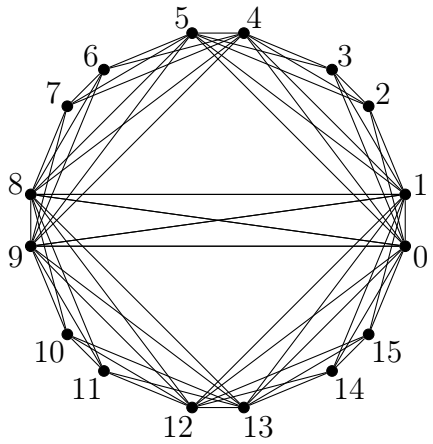

Figure 5-6: A chordal cover of the graph C_{16}^2 obtained as the strong graph product of Γ and K_2, where Γ is the chordal cover of C_8 obtained from Theorem 31 and illustrated in Figure 5-4(left).

where $TM(\theta)$ is the degree d trigonometric moment curve

$$TM(\theta) = \Big(\cos(\theta), \sin(\theta), \cos(2\theta), \sin(2\theta), \ldots, \cos(d\theta), \sin(d\theta)\Big).$$

Cyclic polytopes play an important role in polyhedral combinatorics [102] and satisfy many interesting properties. For example the celebrated *Upper Bound Theorem*, states that for any $2d$ dimensional polytope P with N vertices, $f_i(P) \leq f_i(TC(N, 2d))$ for any $i = 0, \ldots, 2d$, where $f_i(P)$ is the number of i-dimensional faces of a polytope P [102]. Another important property of cyclic polytopes is that they are *neighborly* [45] (recall that a $2d$-dimensional polytope P is called neighborly if any collection of d vertices of P span a face of P).

The results from this section allow us to obtain a positive semidefinite lift of $TC(N, 2d)$ of size $O(d \log(N/d))$ when d divides N. More precisely, if we combine Corollary 3 and Theorem 25D we get that $TC(N, 2d)$ has a Hermitian positive semidefinite lift of size at most $3d \log(N/d)$, proving Theorem 27D from the introduction.

Theorem 27D. *Let N and d be two integers and assume that d divides N. Then the trigonometric cyclic polytope $TC(N, 2d)$ has a Hermitian positive semidefinite lift of size at most $3d \log(N/d)$.*

Comparison with LP lifts One can show that in the regime $N = \Theta(d^2)$ our positive semidefinite lift for $TC(N, 2d)$ is provably smaller than any linear programming lift of $TC(N, 2d)$. Indeed, the following lower bound on the LP extension complexity of k-neighborly polytopes was proved in [40].

Proposition 16. *([40, Proposition 5.16]) If P be a k-neighborly polytope with N vertices then $\mathrm{xc}_{LP}(P) \geq \min(N, (k+1)(k+2)/2)$.*

Since $TC(N, 2d)$ is d-neighborly, if we choose for example $N = d^2$ then the previous proposition asserts that $\mathrm{xc}_{\mathrm{LP}}(TC(d^2, 2d)) \geq \Omega(d^2)$ whereas in this case our positive semidefinite has size $O(d \log d)$. This allows us to prove the following result giving a gap between SDP extension complexity and LP extension complexity.

Corollary 1. *There exists a family $(P_d)_{d\in\mathbb{N}}$ of polytopes where $P_d \subset \mathbb{R}^{2d}$ such that*

$$\frac{\mathrm{xc}_{SDP}(P_d)}{\mathrm{xc}_{LP}(P_d)} = O\left(\frac{\log d}{d}\right).$$

The only nontrivial LP lift for cyclic polytopes that we are aware of is a recent construction by Bogomolov et al. [12] for the cylic polytope

$$C(N, d) = \mathrm{conv}\Big\{(i, i^2, \ldots, i^d) : i = 1, \ldots, N\Big\}$$

of size $(\log N)^{\lfloor d/2 \rfloor}$. Note that this lift has smaller size than the "trivial" vertex lift of $C(N, d)$ only when $d < O((\log N)/(\log \log N))$. Their construction for $d = 2$ is based on the *reflection relations* framework of Kaibel and Pashkovich [62] and the case of general d is then obtained via a tensor product construction, see [12] for details.

5.7 Summary of chapter

- We consider the problem of finding sparse sum-of-squares certificates for nonnegative functions defined on a finite abelian group G that are sparse with respect to the Fourier basis. Using results from the previous chapters, the existence of such certificates translate to (equivariant) SDP lifts of certain moment polytopes (see Equation (5.10)).
- Our main theorem gives a graph-theoretic method to guarantee the existence of such sparse certificates. We apply our main theorem to two specific settings:
- We first show that any nonnegative quadratic function on the hypercube $G = \{-1, 1\}^n$ is a sum of squares of polynomials of degree at most $\lceil n/2 \rceil$ (Theorem 26). This establishes a conjecture of Laurent from 2003 [72] and shows that the Lasserre hierarchy for the cut polytope is exact after $\lceil n/2 \rceil$ levels.
- In the case where $G = \mathbb{Z}_N$ we show that nonnegative functions of degree d on G admit sparse sum-of-squares certificate with support of size $O(d \log(N/d))$ (when d divides N). This allows us to get an explicit sequence of polytopes (trigonometric cyclic polytopes) in increasing dimensions where SDP lifts are vanishingly smaller than LP lifts (Theorem 27 and Corollary 1).

5.8 Proofs

5.8.1 Proof of Theorem 31: chordal cover of the cycle graph

In this appendix we prove Theorem 31 concerning constructing a chordal cover of the cycle graph C_N. Theorem 32 below shows how to construct a chordal cover of the cycle graph C_{N+1} on $N+1$ nodes, by induction. The chordal cover of C_N used to obtain Theorem 31 will then be obtained simply by contracting a certain edge of the chordal cover of C_{N+1} (more details below). We thus start by describing a chordal cover of the $N+1$-cycle.

Theorem 32 (Chordal cover of the cycle graph on $N+1$ vertices)**.** *Let N be an integer greater than or equal 2. Let $k_1 < \cdots < k_l$ be the position of the nonzero digits in the binary expansion of N, i.e., $N = \sum_{i=1}^{l} 2^{k_i}$. Let k be the largest integer such that $2^k < N$ (i.e., $k = k_l - 1$ if N is a power of two and $k = k_l$ otherwise). Then there exists a chordal cover of the cycle graph C_{N+1} on $N+1$ nodes with Fourier support*

$$\mathcal{T} = \{0\} \cup \{\pm 2^i, i = 0, \ldots, k\} \cup \left\{ \sum_{j=1}^{i} 2^{k_j}, i = 1, \ldots, l-1 \right\}. \tag{5.33}$$

Proof. The proof of the theorem is by induction on N. Consider the cycle graph on $N+1$ nodes where nodes are labeled $0, 1, \ldots, N$. To construct a chordal cover of the graph, we first put an edge between nodes 0 and 2^k and another edge between nodes 2^k and N, where 2^k is the largest power of two that is strictly smaller than N. This is depicted in Figure 5-7.

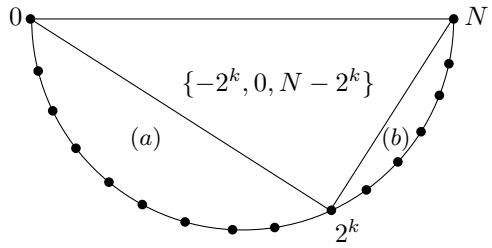

Figure 5-7: Recursive construction of a chordal cover of the cycle $0 \ldots N$ on $N+1$ vertices

Note that the triangle $\{0, 2^k, N\}$ is equivalent, by translation, to $\{-2^k, 0, N-2^k\}$. We now use induction to construct a chordal cover of the two remaining parts of the cycle (denoted (a) and (b) in Figure 5-7).

• For part (a), which is a cycle graph labeled $0 \ldots N'$ with $N' = 2^k$, the induction hypothesis gives us a chordal cover with Fourier support

$$\mathcal{T}_a = \{0\} \cup \{\pm 2^i, i = 0, \ldots, k-1\}. \tag{5.34}$$

• For part (b) of the graph, we use induction on the cycle $2^k \ldots N$ which is, by translation, equivalent to the cycle with labels $0 \ldots N''$ where $N'' = N - 2^k$. We distinguish two cases.

– If $N = 2^{k+1}$, then we have $N'' = 2^k$ and induction gives a chordal cover of (b) with the same Fourier support as for part (a). Thus in this case we get a chordal cover of the full $(N+1)$-cycle with Fourier support

$$\mathcal{T}_a \cup \{-2^k, 0, 2^k\} = \{0\} \cup \{\pm 2^i, i = 0, \ldots, k\}$$

which is what we want.

– Now assume that $N < 2^{k+1}$, which means that the most significant bit of N is at position $k = k_l$. Thus the binary expansion of $N'' = N - 2^k$ is the same as that of N except that the bit at position $k = k_l$ is replaced with a 0. Let k'' be the largest integer such that $2^{k''} < N''$. Using induction we get a chordal cover of the cycle $0 \ldots N''$ with Fourier support

$$\mathcal{T}_b = \{0\} \cup \{\pm 2^i, i = 0, \ldots, k''\} \cup \left\{ \sum_{j=1}^{i} 2^{k_j}, j = 1, \ldots, l-2 \right\}. \tag{5.35}$$

Combining the chordal cover of parts (a) and part (b) we get a chordal cover of the $(N+1)$-cycle with Fourier support

$$\underbrace{\{-2^k, 0, N-2^k\}}_{\text{triangle } \{0, 2^k, N\}} \cup \mathcal{T}_a \cup \mathcal{T}_b.$$

Given the expressions (5.34) and (5.35) for $\mathcal{K}_a$ and $\mathcal{K}_b$, and noting that $k'' \leq k-1$ and that $N - 2^k = \sum_{j=1}^{l-1} 2^{k_j}$, one can check that the chordal cover has Fourier support

$$\mathcal{T} = \{0\} \cup \{\pm 2^i, i = 0, \ldots, k\} \cup \left\{ \sum_{j=1}^{i} 2^{k_j}, i = 1, \ldots, l-1 \right\}.$$

which is exactly what we want.

To complete the proof, it remains to show the base case of the induction. We will show the base cases $N = 2$ and $N = 3$. For $N = 2$, note that the $(N+1)$-cycle is simply a triangle which is already chordal has Fourier support $\{-1, 0, 1\}$. If we evaluate expression (5.33) for $N = 2$ (note that here $k = 0$) we get $\mathcal{T} = \{-1, 0, 1\}$, as needed.

For $N = 3$ (the 4-cycle), we have $k = 1$ and $l = 2$ with $k_1 = 0$ and $k_2 = 1$. Thus expression (5.33) evaluates to $\mathcal{T} = \{0\} \cup \{\pm 1, \pm 2\} \cup \{1\} = \{-2, -1, 0, 1, 2\}$. It is easy to construct a chordal cover of the 4-cycle with such Fourier support (one can even construct one where $\mathcal{T} = \mathcal{K} \cup (-\mathcal{K}) = \{-1, 0, 1\}$). □

Example 17. Figure 5-8 shows the recursive construction for the case $N = 8$. We have indicated in each triangle (3-clique) the associated Fourier support. ◊

Proof of Theorem 31. To prove Theorem 31 for the N-cycle, we use the chordal cover of the $(N+1)$-cycle of Theorem 32 except that we regard nodes 0 and N as the same

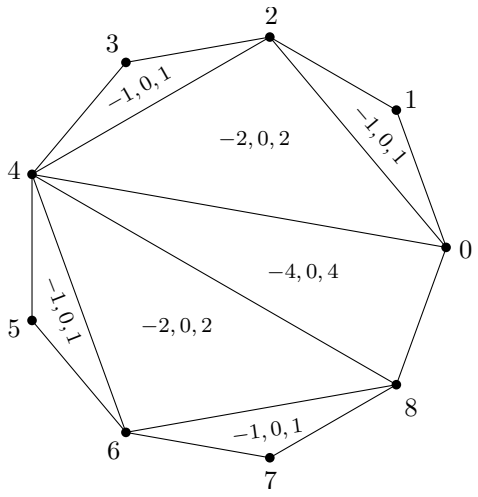

Figure 5-8: Illustration of the recursive chordal cover of the $(N+1)$-cycle for $N = 8$.

nodes (they collapse into a single one). Thus this means that the triangle in Figure 5-7 labeled $\{-2^k, 0, N - 2^k\}$ also collapses and we only have to look at the Fourier support for parts (a) and (b). It is not hard to show that the Fourier support we get is the same as Equation (5.33) except that in the middle term the index i goes from 0 to $k-1$ (instead of from 0 to k), and in the last term the index i goes from 1 to $l-2$ (instead of from 1 to $l-1$). This modification gives exactly the set $\mathcal{T}$ of Equation (5.28). □

Note that there are actually many different ways of constructing chordal covers for the cycle graph, and different constructions will lead to different valid Fourier supports. For instance, for the cycle graph C_N one can actually construct a chordal cover where the size of the Fourier support is related to the logarithm of N base 3. When N is a power of three the Fourier support consists precisely of the powers of 3 that are smaller than N. We omit the precise description of this construction, but Figure 5-9 shows the chordal cover for the 9-cycle and 27-cycle.

5.8.2 Proof of Proposition 12: chordal cover of half-cube graph

Proof of Proposition 12. The proof proceeds as follows. First we define a graph Γ and prove that it is a chordal cover of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$. We then characterize the maximal cliques of Γ. Finally we show that for any maximal clique $\mathcal{C}$ of Γ there is some $S \in 2^{[n]}$ such that $S \triangle \mathcal{C} \subseteq \mathcal{T}$, establishing the stated result. We consider the two cases $\lceil n/2 \rceil$ even and $\lceil n/2 \rceil$ odd separately. We describe the argument in detail in the case where $\lceil n/2 \rceil$ is even, and just sketch the required modifications in the case where $\lceil n/2 \rceil$ is odd.

Assume that $\lceil n/2 \rceil$ is even. Let Γ be the graph with vertex set $2^{[n]}$ such that two vertices S, T are adjacent in Γ if and only if either

- $|S|$ and $|T|$ are both even and $||S| - |T|| \leq 2$ or

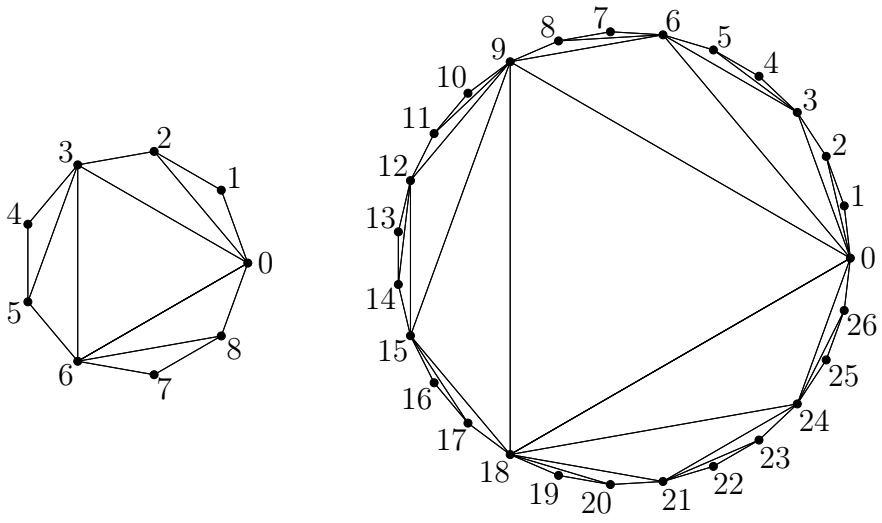

Figure 5-9: Chordal cover of the 9-cycle with Fourier support $\mathcal{T} = \{0, \pm 1, \pm 3\}$ and of the 27-cycle with Fourier support $\mathcal{T} = \{0, \pm 1, \pm 3, \pm 9\}$.

- $|S|$ and $|T|$ are both odd and $||\phi(S)| - |\phi(T)|| \leq 2$.

Note that just like $\mathrm{Cay}(\widehat{G}, \mathcal{S})$, the graph Γ also has two connected components with vertex sets $\mathcal{T}_{\text{even}}$ and $\mathcal{T}_{\text{odd}}$. Furthermore, ϕ (defined in Section 5.5.2) is also an automorphism of Γ that exchanges these two connected components. Observe that if $|S \triangle T| = 2$ (i.e. S and T are adjacent in $\mathrm{Cay}(\widehat{G}, \mathcal{S})$) then both $||S| - |T|| \leq 2$ and $||\phi(S)| - |\phi(T)|| \leq 2$ hold. Hence if S and T are adjacent in $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ they are also adjacent in Γ.

We now show that Γ is a chordal graph. Let the vertices $S_1, S_2, S_3, \ldots, S_k$ form a k-cycle (with $k \geq 4$) in Γ such that each of the $S_i \in \mathcal{T}_{\text{even}}$. Without loss of generality assume that $|S_1| \leq |S_i|$ for $1 \leq i \leq k$. We show that the cycle $S_1, S_2, S_3, \ldots, S_k$ has a chord. If $|S_2| = |S_1|$ then $||S_1| - |S_3|| = ||S_2| - |S_3|| \leq 2$ (since S_2 and S_3 are adjacent) and so there is a chord between S_1 and S_3. Otherwise suppose $|S_2| = |S_1| + 2$. Because S_1 and S_k are adjacent we see that either $|S_k| = |S_1| = |S_2| - 2$ or $|S_k| = |S_1| + 2 = |S_2|$ and so there is a chord between S_2 and S_k. Now suppose $S_1, S_2, S_3, \ldots, S_k$ form a k-cycle (with $k \geq 4$) in Γ such that each of the $S_i \in \mathcal{T}_{\text{odd}}$. Then the image of the cycle under ϕ is a k-cycle in Γ with vertices in $\mathcal{T}_{\text{even}}$ and so it has a chord. Since ϕ is an automorphism of Γ it follows that $S_1, S_2, S_3, \ldots, S_k$ also has a chord. So Γ is a chordal cover of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$.

The subgraphs of Γ induced by the vertex sets $\mathcal{C}_k := \mathcal{T}_k \cup \mathcal{T}_{k+2}$ (for $k = 0, 2, \ldots, 2\lfloor n/2 \rfloor - 2$) and the vertex sets $\phi(\mathcal{C}_k)$ (for $k = 0, 2, \ldots, 2\lfloor n/2 \rfloor - 2$) are cliques in Γ. In fact, these are maximal cliques in Γ. To show that each $\mathcal{C}_k$ is a maximal clique, suppose S is a vertex that is not in $\mathcal{C}_k$. Then either $|S|$ is odd (in which case S is not adjacent to any element of $\mathcal{C}_k$) or $|S| \leq k - 2$ (in which case S is not adjacent to any $T \in \mathcal{T}_{k+2}$) or $|S| \geq k + 4$ (in which case S is not adjacent to any $T \in \mathcal{T}_k$). Hence there is no inclusion-wise larger clique of Γ containing $\mathcal{C}_k$. Since ϕ is an automorphism of Γ it follows that the $\phi(\mathcal{C}_k)$ are also maximal cliques of Γ. Finally, there are no other maximal cliques in Γ because every edge of Γ is contained either in $\mathcal{C}_k$ or $\phi(\mathcal{C}_k)$ for

some $k = 0, 2, \ldots, 2\lfloor n/2 \rfloor - 2$.

It remains to show that for any maximal clique $\mathcal{C}_k$ (for $k = 0, 2, \ldots, 2\lfloor n/2 \rfloor - 2$) of Γ there is $S_k \in 2^{[n]}$ such that $S_k \triangle \mathcal{C}_k \subseteq \mathcal{T}$. This is sufficient to establish that $\mathrm{Cay}(\widehat{G}, \mathcal{S})$ has a chordal cover with Fourier support $\mathcal{T}$ because for the cliques $\phi(\mathcal{C}_k)$ we have that $\phi(S_k) \triangle \phi(\mathcal{C}_k) = S_k \triangle \mathcal{C}_k \subseteq \mathcal{T}$. The following gives valid choices of S_k (for $k = 0, 2, \ldots, 2\lfloor n/2 \rfloor - 2$).

- If $k \leq \lceil n/2 \rceil - 2$ then $\mathcal{C}_k \subseteq \mathcal{T}$ so we can take $S_k = \emptyset$.
- If $k \geq \lceil n/2 \rceil$ and n is even then $n = 2\lceil n/2 \rceil$ and so $n - k - 2 \leq \lceil n/2 \rceil - 2$. Hence $[n] \triangle \mathcal{C}_k = \mathcal{C}_{n-k-2} \subseteq \mathcal{T}$ so we can take $S_k = [n]$.
- If $k \geq \lceil n/2 \rceil$ and n is odd then $n = 2\lceil n/2 \rceil - 1$ and so $n - k + 1 \leq \lceil n/2 \rceil$. Hence

  $$\phi([n]) \triangle \mathcal{C}_k = [n] \triangle \phi(\mathcal{C}_k) \subseteq [n] \triangle (\mathcal{T}_{k-1} \cup \mathcal{T}_{k+1} \cup \mathcal{T}_{k+3}) \subseteq \mathcal{T}_{n-k-3} \cup \mathcal{T}_{n-k-1} \cup \mathcal{T}_{n-k+1} \subseteq \mathcal{T}$$

  so we can take $S_k = \phi([n])$.

This completes the argument in the case where $\lceil n/2 \rceil$ is even.

In the case where $\lceil n/2 \rceil$ is odd we exchange the roles of the odd and even components in the definition of Γ and throughout the argument. More precisely, two vertices S, T are adjacent in Γ if and only if either

- $|S|$ and $|T|$ are both odd and $||S| - |T|| \leq 2$ or
- $|S|$ and $|T|$ are both even and $||\phi(S)| - |\phi(T)|| \leq 2$.

It is still the case that Γ is a chordal cover of $\mathrm{Cay}(\widehat{G}, \mathcal{S})$. Its maximal cliques are now $\mathcal{C}_k := \mathcal{T}_k \cup \mathcal{T}_{k+2}$ for $k = 1, 3, \ldots, 2\lceil n/2 \rceil - 3$ together with the $\phi(\mathcal{C}_k)$. Note that the cliques are now indexed by odd integers. As before, we can choose the S_k (for $k = 1, 3, \ldots, 2\lceil n/2 \rceil - 3$) to be $S_k = \emptyset$ if $k \leq \lceil n/2 \rceil - 2$, $S_k = [n]$ if $k \geq \lceil n/2 \rceil$ and n is even, and $S_k = \phi([n])$ if $k \geq \lceil n/2 \rceil$ and n is odd.

This completes the argument in the case where $\lceil n/2 \rceil$ is odd. $\square$

Chapter 6

Beyond sums of squares: convex proof systems

In Chapter 2 we showed how *certificates of nonnegativity* give us a concrete way to think about lifts of polytopes. We saw in particular that LP and SDP lifts can be understood in terms of producing specific certificates of nonnegativity of the facet inequalities of the polytope.

In this chapter we depart from the specific problem of constructing lifts and we consider the general problem of proving global nonnegativity of an arbitrary real-valued function defined on some domain X. We explore new ways to produce certificates of nonnegativity beyond the LP and SDP framework, using general *convex duality*. The framework we consider unifies existing techniques for certifying nonnegativity, and opens up the possibility of constructing new relaxations for problems that could not be handled using existing technology (e.g., non-polynomial problems). As an illustration we develop in the last section of this chapter a new relaxation for optimization problems involving entropy-like functions (i.e., functions involving terms of the form $x^\alpha \log x$) and we apply it to the problem of computing the so-called *logarithmic Sobolev constant* of finite Markov chains.

6.1 Conic certificates of nonnegativity

The general question of interest in this chapter is to obtain a computationally tractable way of checking whether a given function $f : X \to \mathbb{R}$ is globally nonnegative, i.e., $f(x) \geq 0$ for all $x \in X$. At this point the set X is arbitrary and we will simply assume that f is given as a linear combination of some basis functions $(\phi_i)_{i \in I}$:

$$f = \sum_{i \in I} f_i \phi_i.$$

Conic certificates of nonnegativity We now describe the general form of a *conic certificate of nonnegativity*. Let $K \subset E$ be a given convex cone in some Euclidean space E, and assume we have a map $A : X \to K$ that can be expressed in the basis ϕ_i, i.e.,

$$A(x) = \sum_{i \in I} \phi_i(x) A_i$$

where $A_i \in E$. Such a map A will be called a *lifting map*. If $f : X \to \mathbb{R}$ is a real-valued function on X, we can try to decide the nonnegativity of the function f by solving the following conic feasibility problem:

$$\text{find } B \in K^* \text{ such that } f(x) = \langle A(x), B \rangle \ \forall x \in X. \tag{6.1}$$

Recall that K^* is the dual of K and is defined as

$$K^* = \{y \in E : \langle x, y \rangle \geq 0 \ \forall x \in K\}.$$

As such if $A(x) \in K$ and $B \in K^*$ then $\langle A(x), B \rangle$ is nonnegative, *by definition*. Thus if (6.1) is feasible it provides a certificate of nonnegativity for the function f. Note that the constraint $f(x) = \langle A(x), B \rangle$ for all $x \in X$ can be written in an alternative way in terms of the coefficients of f and A in the basis $(\phi_i)_{i \in I}$, namely as $f_i = \langle A_i, B \rangle$ for all $i \in I$.

Lifting map Of course the main question in coming up with such a proof system is to find a "good" cone K and an associated lifting map $A : X \to K$. We show later how existing certificates of nonnegativity (LP, SOS, geometric programming based method [46], signomials [24], etc.) can be understood as special cases of (6.1) with a well-defined cone K and lifting map A. The framework of conic certificates also allows us to come up with new ways to certify nonnegativity for, say, non-polynomial functions. We show how this can be done in Section 6.6 for a class of *entropy-like* functions.

Connection with lifts of convex sets Certificates of nonnegativity of the form (6.1) have first been proposed by Gouveia, Parrilo, Thomas [50] in the study of K-lifts of convex sets. A convex set C is said to have a K-lift if there exists an affine subspace $L \subset E$ and a linear map π such that $C = \pi(K \cap L)$. Note that LP and SDP lifts

are special cases of K-lifts where K is respectively the nonnegative orthant and the positive semidefinite cone. It was shown in [50] that (under some mild conditions) a convex set C has a K-lift if, and only if, all the valid linear inequalities of C admit a conic certificate of nonnegativity of the form (6.1) (the set X in this case is the set of extreme points of C). This result generalizes Theorem 1 and Theorem 3 on LP and SDP lifts.

Optimization The problem of certifying nonnegativity is in dual relationship with minimization (or maximization) problems. Indeed given a function $f : X \to \mathbb{R}$ the minimum of f on X can be written in the following primal-dual way:

$$f^* = \min_{x \in X} f(x) = \max_{\gamma \in \mathbb{R}} \{\gamma : f - \gamma \text{ is nonnegative on } X\}. \tag{6.2}$$

The equality above simply expresses the fact that the minimum of f on X is equal to the largest global lower bound on f. Thus being able to decide whether for a given $\gamma \in \mathbb{R}$ the function $f - \gamma$ is nonnegative, is equivalent to the problem of minimizing f. Of course the optimization problem (6.2) is hard in general. One way to obtain a relaxation is to replace the constraint that "$f - \gamma$ is nonnegative on X" by "$f - \gamma$ has a conic certificate of nonnegativity of the form (6.1)". This gives the following relaxation of (6.2) whose optimal value is a *lower bound* on f^*:

$$\begin{array}{ll} \text{maximize} & \gamma \\ \text{subject to} & f_0 - \gamma = \langle A_0, B \rangle \\ & f_i = \langle A_i, B \rangle \quad \forall i \in I \setminus \{0\} \\ & B \in K^*. \end{array} \tag{6.3}$$

We assumed here that the function ϕ_0 in the basis $(\phi_i)_{i \in I}$ is the constant function equal to 1 on X. The constraints simply express the fact that $f - \gamma$ has a certificate of nonnegativity of the form $\langle A(x), B \rangle$ where $B \in K^*$. It is useful to look at the (Lagrangian) dual of (6.3), which takes the form:

$$\begin{array}{ll} \text{minimize} & \sum_{i \in I} f_i y_i \\ \text{subject to} & y_0 = 1 \\ & \sum_{i \in I} y_i A_i \in K. \end{array} \tag{6.4}$$

The variables y_i can be interpreted as ϕ_i-moments of a "pseudo-expectation" $\widetilde{E}$ on X, i.e., $y_i = \widetilde{E}[\phi_i(x)]$ (recall the notion of a pseudo-expectation from Sections 2.1.5 and 2.2.5). The constraint $y_0 = 1$ expresses the fact that $\widetilde{E}[1] = 1$ and the constraint $\sum_{i \in I} y_i A_i \in K$ corresponds to $\widetilde{E}_x[A(x)] \in K$ which follows from the fact that $A(x) \in K$ for all $x \in X$ and that K is convex.

Organization In Sections 6.2 to 6.5 we review some of the existing certificates of nonnegativity (LP, sum-of-squares, etc.) in light of the framework described here. In Section 6.6 we propose a new method to certify nonnegativity of certain entropy-like functions and we apply it to the problem of computing logarithmic Sobolev constants

of finite Markov chains.

6.2 LP certificates

LP certificates are probably the simplest way of certifying nonnegativity of functions: assume we have a set of functions $f_1, \ldots, f_N$ on X that we know *a priori* are nonnegative on X. Then given $f : X \to \mathbb{R}$, one way to certify that f is nonnegative is to try to write f as a nonnegative linear combination of $f_1, \ldots, f_N$, i.e., to solve the problem:

$$\text{find } b_1, \ldots, b_N \geq 0 \text{ such that } f = \sum_{i=1}^{N} b_i f_i. \tag{6.5}$$

It is easy to see that this feasibility problem is a special case of (6.1) where the cone K is the nonnegative orthant $\mathbb{R}^N_+$ and the lifting map A is given by:

$$A(x) = \begin{bmatrix} f_1(x) \\ \vdots \\ f_N(x) \end{bmatrix}.$$

Several well-known relaxation methods in optimization are based on LP certificates. Consider for example the problem of optimizing a function f on a set X described using inequalities:

$$X = \{x \in \mathbb{R}^n : g_1(x) \geq 0, \ldots, g_m(x) \geq 0\}.$$

We saw earlier that the minimum of f on X can be expressed in the following primal-dual way:

$$f^* := \min_{x \in X} f(x) = \max_{\gamma \in \mathbb{R}} \{\gamma : f - \gamma \text{ is nonnegative on } X\}. \tag{6.6}$$

One way to certify that $f - \gamma$ is nonnegative on X is to require that there exist coefficients $b_1, \ldots, b_N$ such that $f - \gamma = \sum_{i=1}^{N} b_i g_i$. This yields the following lower bound on f^*:

$$f^* \geq \max_{\gamma \in \mathbb{R}} \left\{ \gamma : f - \gamma = \sum_{i=1}^{N} b_i g_i \text{ where } b_1, \ldots, b_N \geq 0 \right\}. \tag{6.7}$$

This is precisely the so-called *Lagrangian relaxation* of (6.6) and the scalars $b_1, \ldots, b_N$ play the role of the Lagrange variables. In some cases, for example when f and the g_i are affine, we know that we have equality in (6.7), in which case we have *strong duality*.

One can also formulate more general relaxations: it is clear that for any $(\alpha_1, \ldots, \alpha_m) \in \mathbb{N}^m$ the function $x \mapsto g_1(x)^{\alpha_1} \cdots g_m(x)^{\alpha_m}$ is nonnegative on X. Thus one can attempt to show that $f - \gamma$ is nonnegative on X by searching for nonnegative scalars b_α such

that

$$f - \gamma = \sum_{\alpha_1,\ldots,\alpha_m} b_\alpha g_1^{\alpha_1} \cdots g_m^{\alpha_m}. \tag{6.8}$$

Such certificates were considered by Krivine in 1964 [66] where he studied conditions for having a converse (Positivstellensatz) that guarantees the existence of coefficients $b_\alpha \geq 0$ whenever $f - \gamma$ is nonnegative. These certificates also form the basis of the well-known relaxation schemes such as Sherali-Adams [96] or the Handelman hierarchy [58, 75] (see also Section 2.3.1 in Chapter 2 for a related discussion in the context of lifts).

6.3 Sum-of-squares certificates

Another way to certify nonnegativity of functions is using the sum-of-squares method. Let $\psi_1, \ldots, \psi_d$ be a set of functions on X (not necessarily nonnegative) and let V be the subspace of $\mathbb{R}^X$ spanned by these functions. For example if $X = \mathbb{R}$ and $\psi_1(x) = 1, \psi_2(x) = x, \ldots, \psi_d(x) = x^{d-1}$ then V consists of polynomials of degree at most $d-1$. Given a function $f : X \to \mathbb{R}$ one way to certify that f is nonnegative on X is to write f as a sum of squares of functions in V, i.e., to solve the following feasibility problem:

$$\text{find } g_1, \ldots, g_m \in V \text{ such that } f = \sum_{i=1}^m g_i^2. \tag{6.9}$$

Observe that the feasibility problem (6.9) is of the form (6.1) where $K = \mathbf{S}_+^d$ and the lifting map $A_{\boldsymbol{\psi}}$ is given by:

$$A_{\boldsymbol{\psi}}(x) = \boldsymbol{\psi}(x)\boldsymbol{\psi}(x)^T \quad \text{where} \quad \boldsymbol{\psi}(x) = \begin{bmatrix} \psi_1(x) \\ \vdots \\ \psi_d(x) \end{bmatrix}. \tag{6.10}$$

To see why this is the case assume that there exists $B \in (\mathbf{S}_+^d)^* = \mathbf{S}_+^d$ such that

$$f(x) = \langle A_{\boldsymbol{\psi}}(x), B \rangle \quad \forall x \in X. \tag{6.11}$$

Since $B \in \mathbf{S}_+^d$ we can write $B = \sum_{i=1}^m b_i b_i^T$. Thus Equation (6.11) corresponds to

$$f(x) = \left\langle \boldsymbol{\psi}(x)\boldsymbol{\psi}(x)^T, \sum_{i=1}^m b_i b_i^T \right\rangle = \sum_{i=1}^m (b_i^T \boldsymbol{\psi}(x))^2 = \sum_{i=1}^m g_i(x)^2$$

where $g_i = b_i^T \boldsymbol{\psi} \in V$. Conversely it is not hard to show that if (6.9) is feasible, then there exists $B \in \mathbf{S}_+^d$ such that (6.11) is true.

Remark 18. Note that the lifting map (6.10) is rank-one i.e., $\operatorname{rank}(A_{\boldsymbol{\psi}}(x)) = 1$. One can also define higher-rank lifting maps and these correspond to expressing the func-

tion f as a sum of norm squared of vector-valued functions i.e.,

$$f(x) = \sum_{i=1}^{m} \|g_i(x)\|_2^2$$

where $g_i : X \to \mathbb{R}^r$.

We saw in Chapter 2 that sum of squares certificates form the basis of the Lasserre/theta-body relaxations where the function f correspond to the facet inequalities of a polytope P. More generally, sum-of-squares relaxations have been applied in different areas of science and engineering including dynamical systems and control, optimal power flow, quantum information theory, and more. Several results in real algebraic geometry known as *Positivstellensatz* can be used to guarantee the existence of sum-of-squares certificates for positive polynomials. We refer the reader to the books [10, 70] for more details on sum-of-squares relaxations for polynomial optimization problems.

6.4 Geometric programming certificates for homogeneous polynomials

In [46] a method based on *geometric programming* was proposed to certify nonnegativity of polynomials. One appealing aspect of this method is that it results in problems that can be solved faster than the SDPs obtained from the sum-of-squares hierarchy. In this section we show that this method can be cast in the framework of conic certificates of nonnegativity, where the cone K corresponds to a direct product of *power cones*. The main result underlying the method proposed in [46] is the following theorem proved in [39].

Theorem 33 ([39])**.** *Let $a \in \mathbb{N}^n$ with $|a| = 2d$ and consider the homogeneous polynomial in n variables $(x_1, \ldots, x_n)$ of degree $2d$*

$$\sum_{i=1}^{n} b_i x_i^{2d} - \mu x^a \tag{6.12}$$

where $b_i \geq 0$ for all $i = 1, \ldots, n$ and $\mu \in \mathbb{R}$. Assume that $-\mu x^a$ is not a square in $\mathbb{R}[x_1, \ldots, x_n]$ (i.e., either $\mu > 0$ or at least one of the a_i is odd). Then (6.12) is globally nonnegative if and only if

$$\prod_{\substack{i=1 \\ a_i \neq 0}}^{n} \left(\frac{b_i}{a_i/2d} \right)^{a_i/2d} \geq |\mu|. \tag{6.13}$$

Based on this theorem, [46] proposed to certify that a given homogeneous polynomial $f(x)$ of degree $2d$ is nonnegative by trying to express it as a sum of nonnegative polynomials of the form (6.12). The main observation is that searching for such a decomposition can be done using geometric programming.

The cone and the lifting map We now show that such certificates of nonnegativity can be interpreted using a natural lifting map from $\mathbb{R}^n$ to *power cones*. Given $\alpha_i \geq 0$ and $\sum_{i=1}^n \alpha_i = 1$ define the *power cone* $K_{\text{pow},\alpha}$ by:

$$K_{\text{pow},\alpha} = \left\{ (x, z) \in \mathbb{R}^n_+ \times \mathbb{R} : \prod_{i=1}^n x_i^{\alpha_i} \geq |z| \right\}.$$

The cone $K_{\text{pow},\alpha}$ is a well-known closed convex cone, see e.g., [25] for properties and applications of this cone. For any $a \in \mathbb{N}^n$ with $|a| := \sum_{i=1}^n a_i = 2d$ we can consider the following lifting map $A_{\text{pow},a} : \mathbb{R}^n \to K_{\text{pow},a/2d}$:

$$A_{\text{pow},a}(x) = \begin{bmatrix} x_1^{2d} \\ \vdots \\ x_n^{2d} \\ x^a \end{bmatrix} \in K_{\text{pow},a/2d}.$$

Note that $A_{\text{pow},a}(x) \in K_{\text{pow},a/2d}$ because $\prod_{i=1}^n (x_i^{2d})^{a_i/2d} = \prod_{i=1}^n |x_i|^{a_i} = |x^a|$. The following proposition shows that one can characterize the set of nonnegative polynomials of the form (6.12) in terms of the dual cone $(K_{\text{pow},a/2d})^*$.

Proposition 17. *Let $f(x) = \sum_{i=1}^n b_i x_i^{2d} - \mu x^a$ and assume that $b_i \geq 0$ for all $i = 1, \ldots, n$ and that $-\mu x^a$ is not a square in $\mathbb{R}[x_1, \ldots, x_n]$. Then f is globally nonnegative if and only if there exists $B \in (K_{pow,a/2d})^*$ such that $f(x) = \langle A_{pow,a}(x), B \rangle$.*

Proof. We can prove this proposition in two ways. The first way is to get an analytical expression of the dual of $K_{\text{pow},a/2d}$ and show that the conditions $B \in (K_{\text{pow},a/2d})^*$ and $f(x) = \langle A_{\text{pow},a}(x), B \rangle$ coincide with the conditions (6.13). We prefer however to prove the proposition directly using simply the definition of conic duality.

Let $f(x) = \sum_{i=1}^n b_i x_i^{2d} - \mu x^a$ and assume that $b_i \geq 0$ for all $i = 1, \ldots, n$ and that $-\mu x^a$ is not a square in $\mathbb{R}[x_1, \ldots, x_n]$. First, note that given $B \in \mathbb{R}^{n+1}$ we have, by matching coefficients,

$$f(x) = \langle A_{\text{pow},a}(x), B \rangle \quad \iff \quad B = (b_1, \ldots, b_n, -\mu).$$

(we assume here that $a_i \neq 2d$ for all $i = 1, \ldots, n$; the case where there is i_0 such that $a_{i_0} = 2d$ and $a_i = 0$ for $i \neq i_0$ can be easily treated separately and we omit the details here). Thus to prove the proposition we need to show that $f(x)$ is globally nonnegative if and only if $(b_1, \ldots, b_n, -\mu) \in (K_{\text{pow},a/2d})^*$.

We are first going to treat the case $\mu \geq 0$ for simplicity. The case $\mu \leq 0$ will then

follow. We have the following equivalences:

$$\begin{aligned}
\sum_{i=1}^{n} b_i x_i^{2d} - \mu x^a \geq 0 \quad \forall x \in \mathbb{R}^n \quad &\overset{(a)}{\iff} \quad \sum_{i=1}^{n} b_i x_i^{2d} - \mu x^a \geq 0 \quad \forall x \in \mathbb{R}^n_+ \\
&\overset{(b)}{\iff} \quad \sum_{i=1}^{n} b_i y_i - \mu y^{a/2d} \geq 0 \quad \forall y \in \mathbb{R}^n_+ \\
&\overset{(c)}{\iff} \quad \sum_{i=1}^{n} b_i y_i - \mu z \geq 0 \quad \forall (y,z) \in \mathbb{R}^n_+ \times \mathbb{R} : y^{a/2d} \geq |z| \\
&\overset{(d)}{\iff} \quad (b_1, \ldots, b_n, -\mu) \in (K_{\text{pow},\alpha})^*
\end{aligned}$$

In (a), for the implication $\Leftarrow$ we use the assumption that $\mu \geq 0$ which implies that for any $x \in \mathbb{R}^n$ we have $\sum_{i=1}^n b_i x_i^{2d} - \mu x^a \geq \sum_{i=1}^n b_i x_i^{2d} - \mu |x|^a \geq 0$. For (b) we simply used the change of variables $y_i = x_i^{2d}$ which implies that $y_i^{a_i/2d} = x_i$ for $x \geq 0$. Step (c) is trivial using the fact that $\mu \geq 0$ and step (d) is simply the definition of $(K_{\text{pow},\alpha})^*$. Thus this completes the proof in the case $\mu \geq 0$.

Now if $\mu \leq 0$ since we assumed that $-\mu x^a$ is not a square, this means that at least one of the a_i, say a_{i_0} is odd. By doing a change of variables $\widetilde{x}_{i_0} = -x_{i_0}$ we see that the following is true:

$$\sum_{i=1}^{n} b_i x_i^{2d} - \mu x^a \geq 0 \quad \forall x \in \mathbb{R}^n \iff \sum_{i=1}^{n} b_i x_i^{2d} + \mu x^a \geq 0 \quad \forall x \in \mathbb{R}^n.$$

By the reasoning above the latter is equivalent to having $(b_1, \ldots, b_n, \mu) \in (K_{\text{pow},\alpha})^*$. By definition of $K_{\text{pow},\alpha}$ it is easy to see that in turn this is equivalent to having $(b_1, \ldots, b_n, -\mu) \in (K_{\text{pow},\alpha})^*$ which proves the claim in the case $\mu \leq 0$.

□

Remark 19. Note that one can get a closed-form expression for the dual of the power cone (see for example [25]):

$$(K_{\text{pow},\alpha})^* = \left\{ (x^*, z^*) \in \mathbb{R}^n_+ \times \mathbb{R} : \prod_{\substack{i=1 \\ \alpha_i \neq 0}}^{n} \left(\frac{x_i^*}{\alpha_i} \right)^{\alpha_i} \geq |z^*| \right\}. \tag{6.14}$$

Thus we can verify that condition (6.13) of Theorem 33 can be written as $(b_1, \ldots, b_n, -\mu) \in (K_{\text{pow},a/2d})^*$.

Given a homogeneous polynomial f of degree $2d$ in n variables, the approach of [46] to certify nonnegativity of f is to try to express f as a sum of nonnegative functions of the form (6.12). Let $\text{supp}(f) \subset \mathbb{N}^n$ be the support of f, i.e.,

$$\text{supp}(f) = \{a \in \mathbb{N}^n : f_a \neq 0\}$$

where f_a is the coefficient of the monomial x^a in the expansion of f in the monomial

basis. Let $\Delta(f)$ be the set of monomials in f that are not squares

$$\begin{aligned}\Delta(f) &= \{a \in \text{supp}(f) \text{ and } f_a x^a \text{ is not a square in } \mathbb{R}[x_1, \ldots, x_n]\} \\ &= \{a \in \text{supp}(f) : f_a < 0 \text{ or at least one of the } a_i \text{ is odd}\}.\end{aligned}$$

If $a \in \text{supp}(f) \setminus \Delta(f)$ then a_i is even for all i and we can consider the following map into $\mathbb{R}_+$,

$$A_{\text{sq},a} : x \in \mathbb{R}^n \mapsto x^a \in \mathbb{R}_+.$$

Consider the lifting map A_f from $\mathbb{R}^n$ to the Cartesian product of cones

$$K = \prod_{a \in \Delta(f)} K_{\text{pow},a/2d} \times \prod_{a \in \text{supp}(f) \setminus \Delta(f)} \mathbb{R}_+$$

obtained by stacking together the maps $A_{\text{pow},a}$ for $a \in \Delta(f)$ and the $A_{\text{sq},a}$ for $a \in \text{supp}(f) \setminus \Delta(f)$. To certify that f is globally nonnegative we can search for $B \in K^*$ such that:

$$f(x) = \langle A_f(x), B \rangle.$$

This corresponds exactly to the geometric programming approach of Ghasemi and Marshall [46, Theorem 2.3].

6.5 Signomials

Consider the set of functions defined on $X = \mathbb{R}^n$ of the the form

$$\sum_{i=1}^{N} c_i \exp(\alpha_i^T x) \tag{6.15}$$

where $\alpha_1 \in \mathbb{R}^n, \ldots, \alpha_N \in \mathbb{R}^n$, and $c_1, \ldots, c_N \in \mathbb{R}$. Such functions are known as *signomial functions*. In [24] a method based on convex optimization was proposed to certify nonnegativity of signomial functions. In this section we show that the method of [24] can be understood in the framework described here with a natural lifting map from $\mathbb{R}^n$ to a certain product of exponential cones.

SAGE certificates We first recall the method proposed in [24] to certify nonnegativity of signomial functions. The main result underlying the method of [24] is stated in Proposition 18 below and shows that one can *exactly* certify nonnegativity of (6.15) when at most one of the coefficients c_i is negative. For the statement of the proposition, define

$$D_{\text{KL}}(u \| v) := \sum_{j=1}^{l} u_j \log(u_j / v_j)$$

to be the *relative entropy* function defined for $(u, v) \in \mathbb{R}^l_{++} \times \mathbb{R}^l_{++}$. The function D_{KL} is well-known to be convex. The following proposition was proved in [24]:

Proposition 18 ([24]). *Consider a signomial function of the form*

$$\sum_{j=1}^{l} c_j \exp(\alpha_j^T x) + \overline{c} \exp(\overline{\alpha}^T x) \tag{6.16}$$

where $c_j \geq 0$ for $j = 1, \ldots, l$ and $\overline{c} \in \mathbb{R}$. Then (6.16) is globally nonnegative if and only if there exists $\nu \in \mathbb{R}^l_+$ such that

$$D_{KL}(\nu \| e \cdot c) \leq \overline{c} \quad \textit{and} \quad \sum_{j=1}^{l} \nu_j (\overline{\alpha} - \alpha_j) = 0 \tag{6.17}$$

where e is the constant $e = \exp(1)$.

Note that since D_{KL} is a convex function the conditions (6.17) are convex in $(\nu, c, \overline{c})$. Given Proposition 18, the approach proposed in [24] to certify that a signomial function (6.15) is nonnegative is to try to express it as a sum of nonnegative signomial functions where each one has at most one negative coefficient. Since the conditions (6.17) are convex one can search over such decompositions efficiently using convex optimization. More precisely, given $\alpha_1, \ldots, \alpha_N \in \mathbb{R}^n$, define $\mathrm{AGE}(\alpha_i; \alpha_{-i})$ to be the set of signomial functions of the form (6.15) that are nonnegative and where $c_j \geq 0$ for all $j \neq i$ (here the notation α_{-i} refers to $(\alpha_j)_{j \neq i}$). Then define $\mathrm{SAGE}(\alpha_1, \ldots, \alpha_N)$ to be the set of functions that can be expressed as a sum of functions each in $\mathrm{AGE}(\alpha_i; \alpha_{-i})$ for $i = 1, \ldots, N$, i.e.,:

$$\mathrm{SAGE}(\alpha_1, \ldots, \alpha_N) = \mathrm{AGE}(\alpha_1; \alpha_{-1}) + \cdots + \mathrm{AGE}(\alpha_N; \alpha_{-N}). \tag{6.18}$$

In [24] a nonnegative signomial function with at most one negative coefficient is called a "AM-GM exponential", hence the notation AGE. Checking that a function is in $\mathrm{SAGE}(\alpha_1, \ldots, \alpha_N)$ can be done using convex optimization, and more precisely using *relative entropy optimization* [23].

The cone and the lifting map We now give an interpretation of SAGE certificates (6.18) using a naturally defined lifting map from $X = \mathbb{R}^n$ to a certain product of exponential cones. Let $K^{(l)}_{\mathrm{AGE}}$ be the cone in $\mathbb{R}^{2l+1}$ defined as:

$$K^{(l)}_{\mathrm{AGE}} = \mathrm{cl}\left\{(t, y, z) \in \mathbb{R}^l \times \mathbb{R}_{++} \times \mathbb{R}^l_+ : y e^{t_j/y} \leq z_j \; \forall j = 1, \ldots, l\right\}.$$

Note that $K^{(l)}_{\mathrm{AGE}}$ is related to the well-known *exponential cone*:

$$K_{\exp} = \mathrm{cl}\left\{(t, y, z) \in \mathbb{R} \times \mathbb{R}_{++} \times \mathbb{R}_{++} : y e^{t/y} \leq z\right\}. \tag{6.19}$$

Indeed $K^{(l)}_{\mathrm{AGE}}$ is the intersection of the Cartesian product $(K_{\exp})^l = K_{\exp} \times \cdots \times K_{\exp}$ with the subspace $\{(t, y, z) \in \mathbb{R}^{3l} : y_1 = \cdots = y_l\}$.

Given $\alpha_1, \ldots, \alpha_l, \overline{\alpha} \in \mathbb{R}^n$, consider the map $A_{\overline{\alpha};\alpha} : \mathbb{R}^n \to K^{(l)}_{\text{AGE}}$:

$$A_{\overline{\alpha};\alpha}(x) = \begin{bmatrix} \dfrac{\left[\exp(\overline{\alpha}^T x) \cdot (\alpha_j - \overline{\alpha})^T x\right]_{j=1,\ldots,l}}{\exp(\overline{\alpha}^T x)} \\ \hline \left[\exp(\alpha_j^T x)\right]_{j=1,\ldots,l} \end{bmatrix}. \tag{6.20}$$

It is a straightforward calculation to verify that $A_{\overline{\alpha};\alpha}(x) \in K^{(l)}_{\text{AGE}}$. The next proposition shows that one can characterize the set of nonnegative signomial functions with at most one negative coefficient in terms of $A_{\overline{\alpha};\alpha}$.

Proposition 19. *Consider the signomial f defined in (6.16) where $c_j \geq 0$ for all $j = 1, \ldots, l$ and $\overline{c} \in \mathbb{R}$. Then f is globally nonnegative if and only if there exists $B \in (K^{(l)}_{AGE})^*$ such that $f(x) = \langle A_{\overline{\alpha};\alpha}(x), B\rangle$.*

Proof. Similarly to Proposition 17, one can prove this proposition in two ways. One way is to compute an analytical expression for $(K^{(l)}_{\text{AGE}})^*$ and show that the conditions $B \in (K^{(l)}_{\text{AGE}})^*$ and $f(x) = \langle A_{\overline{\alpha};\alpha}(x), B\rangle$ coincide with (6.17). We give here another proof that simply uses the definition of duality for convex cones and does not involve computing the dual $(K^{(l)}_{\text{AGE}})^*$.

- We first simplify the condition $f(x) = \langle A_{\overline{\alpha},\alpha}(x), B\rangle$ and $B \in (K^{(l)}_{\text{AGE}})^*$. Observe that if $B = (t^*, y^*, z^*) \in \mathbb{R}^l \times \mathbb{R} \times \mathbb{R}^l$ then by matching coefficients we have:

$$f(x) = \langle A_{\overline{\alpha};\alpha}(x), B\rangle \iff \begin{cases} z_j^* = c_j \quad \forall j = 1, \ldots, l \\ y^* = \overline{c} \\ \sum_{j=1}^l t_j^*(\alpha_j - \overline{\alpha}) = 0. \end{cases} \tag{6.21}$$

  For convenience let us introduce the matrix $M \in \mathbb{R}^{l \times n}$ given by:

$$M = \begin{bmatrix} (\alpha_1 - \overline{\alpha})^T \\ \vdots \\ (\alpha_l - \overline{\alpha})^T \end{bmatrix} \in \mathbb{R}^{l \times n}.$$

  The last condition on t^* in the right-hand side of (6.21) is equivalent to $t^* \in \text{Ker}(M^T) = \text{Im}(M)^\perp$. Thus from (6.21) we have the equivalence:

$$\begin{array}{c} \exists B \in (K^{(l)}_{\text{AGE}})^* \text{ s.t. } f(x) = \langle A_{\overline{\alpha};\alpha}(x), B\rangle \\ \Updownarrow \\ \exists t^* \in \text{Im}(M)^\perp \text{ s.t. } (t^*, \overline{c}, c) \in (K^{(l)}_{\text{AGE}})^*. \end{array} \tag{6.22}$$

- Second, observe that nonnegativity of the signomial $f(x)$ can be equivalently

written as:

$$\begin{aligned}
&\sum_{j=1}^{l} c_j \exp(\alpha_j^T x) + \overline{c}\exp(\overline{\alpha}^T x) \geq 0 \quad \forall x \in \mathbb{R}^n \\
\iff &\sum_{j=1}^{l} c_j \exp((\alpha_j - \overline{\alpha})^T x) + \overline{c} \geq 0 \quad \forall x \in \mathbb{R}^n \\
\iff &\sum_{j=1}^{l} c_j \exp(t_j) + \overline{c} \geq 0 \quad \forall t \in \mathrm{Im}(M) \\
\iff &\sum_{j=1}^{l} c_j z_j + \overline{c} \geq 0 \quad \forall (t,z) \in \mathrm{Im}(M) \times \mathbb{R}^l_+ \text{ s.t. } e^{t_j} \leq z_j, j = 1, \ldots, l \\
\iff &\sum_{j=1}^{l} c_j z_j + \overline{c} y \geq 0 \quad \forall (t,y,z) \in \mathrm{Im}(M) \times \mathbb{R}_{++} \times \mathbb{R}^l_+ \\
&\qquad\qquad\qquad\qquad\qquad\qquad \text{s.t. } y e^{t_j/y} \leq z_j, j = 1, \ldots, l.
\end{aligned}$$

The last condition is, by definition of duality, equivalent to saying that $(0, \overline{c}, c) \in (K^{(l)}_{\mathrm{AGE}} \cap L)^*$ where $L \subset \mathbb{R}^{2l+1}$ is the subspace $L = \mathrm{Im}(M) \times \mathbb{R} \times \mathbb{R}^l$. Using well-known properties of conic duality we have $(K^{(l)}_{\mathrm{AGE}} \cap L)^* = (K^{(l)}_{\mathrm{AGE}})^* + L^\perp$ and $L^\perp = \mathrm{Im}(M)^\perp \times \{0\} \times \{0\}$. Thus this means that

$$(0, \overline{c}, c) \in (K^{(l)}_{\mathrm{AGE}} \cap L)^* \iff \exists t^* \in \mathrm{Im}(M)^\perp \text{ s.t. } (t^*, \overline{c}, c) \in (K^{(l)}_{\mathrm{AGE}})^*. \tag{6.23}$$

This completes the proof since it is the same as condition (6.22). □

Consider now a general signomial function of the form (6.15). For each $i = 1, \ldots, N$, we can consider the lifting map $A_{\alpha_i, \alpha_{-i}} : \mathbb{R}^n \to K^{(N-1)}_{\mathrm{AGE}}$ where $\alpha_{-i} = (\alpha_j)_{j \neq i}$. Define

$$A_{\mathrm{SAGE},\alpha} : \mathbb{R}^n \to K^{(N-1)}_{\mathrm{AGE}} \times \cdots \times K^{(N-1)}_{\mathrm{AGE}} = \left(K^{(N-1)}_{\mathrm{AGE}}\right)^N$$

obtained by stacking together all the $A_{\alpha_i, \alpha_{-i}}$, i.e.,

$$A_{\mathrm{SAGE},\alpha}(x) = \left[A_{\alpha_i, \alpha_{-i}}(x)\right]_{i=1,\ldots,N} \in \left(K^{(N-1)}_{\mathrm{AGE}}\right)^N.$$

Then it is easy to see that a signomial f is in $\mathrm{SAGE}(\alpha_1, \ldots, \alpha_N)$ if and only if there exists $B \in \left(\left(K^{(N-1)}_{\mathrm{AGE}}\right)^N\right)^*$ such that $f(x) = \langle A_{\mathrm{SAGE},\alpha}(x), B \rangle$.

6.6 New certificates for entropy-like functions and applications

In this section we introduce new certificates of nonnegativity for a class of entropy-like functions, i.e., functions of the form:

$$p_0(x) + \sum_{i=1}^{n} p_i(x) \log x_i \tag{6.24}$$

where $p_0, p_1, \ldots, p_n$ are polynomials and $x \in \mathbb{R}^n_{++}$. We first propose a proof system relying on the concavity of log. We then see how to improve this proof system by exploiting the *matrix concavity* of the logarithm. Finally we discuss an application of our method to the problem of computing the *log-Sobolev constant* of a given finite Markov chain.

6.6.1 Certificates based on classical relative entropy

In this section we see how to exploit the concavity of the logarithm function to derive a proof system for functions of the form (6.24).

Concavity of log Since the function log is concave one can easily characterize all the nonnegative univariate functions of the form

$$\alpha + \beta x - \log x \tag{6.25}$$

where $\alpha, \beta \in \mathbb{R}$. Indeed nonnegative functions of the form (6.25) correspond to affine functions that lie above the curve of log and are encoded in the Fenchel conjugate of log. From (6.25) one can generate nonnegative functions on $\mathbb{R}^n_{++}$ as follows:

> If f is a univariate nonnegative function of the form (6.25) (or a nonnegative multiple of it) then $x \in \mathbb{R}^n_{++} \mapsto x^r f(x^{s-r})$ is nonnegative on $\mathbb{R}^n_{++}$ where $r, s \in \mathbb{N}^n$. (6.26)

In (6.26) the notation x^r corresponds to $\prod_{i=1}^n x_i^{r_i}$. For example if $f(x) = x - 1 - \log(x)$ (which is nonnegative and corresponds to $\alpha = -1$ and $\beta = 1$ in (6.25)) and if we take $r = (2, 1) \in \mathbb{N}^2$ and $s = (1, 0) \in \mathbb{N}^2$ we get that the bivariate function of x_1, x_2

$$x_1^2 x_2 (x_1^{-1} x_2^{-1} - 1 - \log(x_1^{-1} x_2^{-1})) = x_1 - x_1^2 x_2 + x_1^2 x_2 \log(x_1) + x_1^2 x_2 \log(x_2)$$

is nonnegative on $\mathbb{R}^2_{++}$. By varying $(r, s) \in \mathbb{N}^n \times \mathbb{N}^n$ we get from (6.26) different classes of nonnegative functions that are all of the form (6.24). Thus in order to certify that an entropy-like function is nonnegative, we can try to express it as a sum of functions obtained from (6.26) with different values of $(r, s) \in \mathbb{N}^n \times \mathbb{N}^n$. This defines a proof system for entropy-like functions of the form (6.24), and we will see that it can be cast in the framework considered in this chapter, by defining an appropriate lifting map from $\mathbb{R}^n_{++}$ to the *relative entropy cone*.

Relative entropy cone and lifting map The relative entropy cone, denoted K_{re}, is the epigraph of the relative entropy function $(u, v) \mapsto u \log(u/v)$:

$$K_{\text{re}} = \text{cl}\left\{(u, v, w) \in \mathbb{R}_{++} \times \mathbb{R}_{++} \times \mathbb{R} : u \log\left(\frac{u}{v}\right) \leq w\right\}. \tag{6.27}$$

The cone K_{re} is related to the exponential cone K_{exp} that we considered in a previous section (cf. Equation (6.19)). In fact one can show that:

$$(u, v, w) \in K_{\text{re}} \Leftrightarrow (x = -w, y = u, z = v) \in K_{\text{exp}}.$$

We are now going to define a lifting map from $\mathbb{R}^n_{++}$ to K_{re} which allows us to capture the certificates of nonnegativity described in (6.26). Given $(r, s) \in \mathbb{N}^n \times \mathbb{N}^n$, consider the map $A_{r,s} : \mathbb{R}^n_{++} \to K_{\text{re}}$ defined by:

$$A_{r,s}(x) = \begin{bmatrix} x^r \\ x^s \\ x^r (r-s)^T \log x \end{bmatrix}. \tag{6.28}$$

In (6.28) the notation $\log x$ for $x = (x_1, \ldots, x_n)$ corresponds to the vector $(\log x_1, \ldots, \log x_n)$ and the term $(r-s)^T \log x$ thus corresponds to $\sum_{i=1}^n (r_i - s_i) \log x_i$. It is straightforward to verify that $A_{r,s}(x) \in K_{\text{re}}$ for all $x \in \mathbb{R}^n_{++}$. Indeed we have:

$$x^r \log\left(\frac{x^r}{x^s}\right) = x^r \log\left(x^{r-s}\right) = x^r \sum_{i=1}^n (r_i - s_i) \log x_i = x^r (r-s)^T \log x.$$

The next proposition shows that nonnegative functions obtained from the lifting map (6.28) are precisely those obtained from (6.26).

Proposition 20. *Let $r, s \in \mathbb{N}^n$ fixed. The nonnegative functions obtained according to (6.26) are precisely those of the form $x \mapsto \langle A_{r,s}(x), B\rangle$ where $B \in K^*_{re}$.*

Proof. Observe that since $A_{r,s}(x) = x^r A_{0,1}(x^{s-r})$, the functions of the form $x \mapsto \langle A_{r,s}(x), B\rangle$ are exactly those of the form $x \mapsto x^r f(x^{s-r})$ where f is a univariate function $f(x) = \langle A_{0,1}(x), B\rangle$. Thus to prove the proposition we just need to consider the case $r = 0, s = 1$ and show that the nonnegative univariate functions of the form $\alpha + \beta x - \gamma \log x$ (where $\gamma \geq 0$) are exactly those that can be written as $\langle A_{0,1}(x), B\rangle$ where $B \in K^*_{\text{re}}$. Since for any $B = (\alpha, \beta, \gamma) \in \mathbb{R}^3$ we have

$$\langle A_{0,1}(x), B\rangle = \alpha + \beta x - \gamma \log x$$

what we just need to show is that $\alpha + \beta x - \gamma \log x$ is nonnegative on $\mathbb{R}_{++}$ if and only if $(\alpha, \beta, \gamma) \in K^*_{\text{re}}$. This follows directly from the definition of K^*_{re}, indeed:

$$\begin{aligned} (\alpha, \beta, \gamma) \in K^*_{\text{re}} &\iff \alpha u + \beta v + \gamma w \geq 0 \quad \forall (u, v, w) : u \log(u/v) \leq w \\ &\overset{(a)}{\iff} \alpha u + \beta v + \gamma u \log(u/v) \geq 0 \quad \forall u, v > 0 \\ &\overset{(b)}{\iff} \alpha + \beta x - \gamma \log(x) \geq 0 \quad \forall x > 0 \end{aligned} \tag{6.29}$$

where in (a) we used the fact that $\gamma \geq 0$ and in (b) we used the change of variables $x = v/u$. □

Examples We now give two simple examples to illustrate the proof system we just defined.

Example 18 (Univariate example). Let $f(x) = x^2 - 3x + 2 + x\log x$. From the nonnegative function $x - 1 - \log x$ and using different choices of (r, s), we get according to (6.26) the following inequalities which are valid on $\mathbb{R}_{++}$:

$$\begin{aligned} (r = 1, s = 0) \qquad & -x + 1 + x\log(x) \geq 0 \\ (r = 1, s = 2) \qquad & x^2 - x - x\log(x) \geq 0 \end{aligned}$$

By multiplying the first inequality by 2 and adding it to the second inequality we get that $f(x) = x^2 - 3x + 2 + x\log x \geq 0$. ◊

Example 19 (Motzkin polynomial). Let

$$M(x, y) = x^4y^2 + x^2y^4 + 1 - 3x^2y^2$$

be the dehomogenized Motzkin polynomial which is globally nonnegative on $\mathbb{R}^2$. Note that $M(x, y) = P(x^2, y^2)$ where

$$P(x, y) = x^2y + xy^2 + 1 - 3xy.$$

Since M is globally nonnegative, P is nonnegative on $\mathbb{R}^2_+$. The certificates of nonnegativity we just defined allow us to prove nonnegativity of P on $\mathbb{R}^2_+$. Indeed from the inequality $x - 1 - \log(x) \geq 0$ and using different choices of r, s we get from (6.26) the following inequalities which are valid on $\mathbb{R}^2_{++}$:

$$\begin{aligned} (r = (1,1), s = (0,0)) \qquad & -xy + 1 + xy(\log(x) + \log(y)) \geq 0 \\ (r = (1,1), s = (2,1)) \qquad & -xy + x^2y + xy(-\log(x) + \log(y)) \geq 0 \\ (r = (1,1), s = (1,2)) \qquad & -xy + xy^2 + xy(\log(x) - \log(y)) \geq 0 \end{aligned}$$

Summing these inequalities gives $P(x, y) \geq 0$ for all $x, y > 0$. ◊

Operator concavity of logarithm The proof system we just defined exploits the concavity of the logarithm. It turns out however that the logarithm function is concave in a much stronger sense, in fact it is *operator concave*. This means that for any integer $n \geq 1$ and any two positive definite matrices $X, Y \in \mathbf{S}^n_{++}$ and any $\lambda \in [0, 1]$ the following matrix inequality holds:

$$\log\left(\lambda X + (1 - \lambda)Y\right) \succeq \lambda \log(X) + (1 - \lambda)\log(Y).$$

The operator concavity of logarithm can be exploited to define a stronger proof system for entropy-like functions. In the next section we first review some basic results concerning operator concavity before proceeding to the definition of such a proof system.

6.6.2 Operator convexity

In this section we review some results dealing with convexity of matrix-valued maps. Good references include [8, 19]. If $f : I \to \mathbb{R}$ is a function defined on an interval I of $\mathbb{R}$ and if A is a symmetric matrix with eigenvalues in I then we define $f(A)$ via the spectral decomposition of A:

$$f(A) = \sum_i f(\lambda_i) P_i$$

where

$$A = \sum_i \lambda_i P_i$$

is the spectral decomposition of A (P_i is the orthogonal projection on the eigenspace associated to λ_i).

Definition 19 (Operator convex functions)**.** Let $f : I \to \mathbb{R}$ where I is an interval of $\mathbb{R}$. We say that f is *operator convex* if for any integer $d \geq 1$ and any $X, Y \in \mathbf{S}^d$ with eigenvalues in I and $\lambda \in [0, 1]$ we have:

$$f(\lambda X + (1 - \lambda)Y) \preceq \lambda f(X) + (1 - \lambda) f(Y). \tag{6.30}$$

The definition of *operator concave* is similar except that the sign $\preceq$ is replaced by $\succeq$ in (6.30).

Clearly if f is operator convex then it must be convex in the usual sense (this is obtained by taking $d = 1$ in the definition). However not all convex functions are operator convex. For example the function $f(x) = x^4$ is convex on $\mathbb{R}$ but not operator convex. One can check for example that the inequality (6.30) is violated with $X = \left[\begin{smallmatrix} 1 & 0 \\ 0 & 2 \end{smallmatrix}\right], Y = \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right], \lambda = 1/2$. An important result however due to Löwner is the following.

Theorem 34 (Löwner)**.** *For any $t \in [0, 1]$ the function $f(x) = x^t$ is operator concave.*

One consequence of this theorem is that the logarithm function is operator concave. Indeed one can easily verify that for any $X \succ 0$

$$\log X = \lim_{t \to 0} \frac{1}{t}(X^t - I). \tag{6.31}$$

Since for any $t \in [0, 1]$ the map $\frac{1}{t}(x^t - 1)$ is operator concave, it follows that $\log x$ is also operator concave.

Matrix perspective It is well-known in convex analysis that if f is a convex function whose domain is a convex cone, then the perspective of f defined by $g(x, y) = yf(x/y)$ on $\operatorname{dom} f \times \mathbb{R}_{++}$ is also convex. A similar result can be shown to hold in the matrix case where the scaling parameter y is a positive definite matrix (instead of being a scalar).

Theorem 35. *Let $f : \mathbb{R}_{++} \to \mathbb{R}$ and assume that f is operator convex. For any integer $d \geq 1$, define g by*

$$g(X, Y) = Y^{1/2} f\left(Y^{-1/2} X Y^{-1/2}\right) Y^{1/2}$$

where $X \in \mathbf{S}^d_{++}$ and $Y \in \mathbf{S}^d_{++}$. Then g is convex in the Löwner ordering, namely for any (X_1, Y_1), (X_2, Y_2) and $\lambda \in [0, 1]$ we have

$$g\left(\lambda X_1 + (1-\lambda)X_2, \lambda Y_1 + (1-\lambda)Y_2\right) \preceq \lambda g(X_1, Y_1) + (1-\lambda) g(X_2, Y_2).$$

Proof. See [98, Theorem 8.6.2]. □

Remark 20. An equivalent way of formulating the convexity of g in the Löwner ordering, is to say that its *matrix epigraph*, defined by

$$\{(X, Y, T) \in \mathbf{S}^d_{++} \times \mathbf{S}^d_{++} \times \mathbf{S}^d : g(X, Y) \preceq T\}$$

is convex.

6.6.3 Certificates based on matrix relative entropy

Define Ψ to be the matrix perspective function of log:

$$\Psi(A, B) = A^{1/2} \log\left(A^{-1/2} B A^{-1/2}\right) A^{1/2} \quad \forall (A, B) \in \mathbf{S}^n_{++} \times \mathbf{S}^n_{++}. \tag{6.32}$$

Since log is operator concave, Theorem 35 asserts that Ψ is also concave in the Löwner ordering. The function Ψ was considered by Fujii and Kamei [43] as a matrix-valued relative entropy. Some properties of this function are studied in [42] and [44].

Cone and lifting map Since Ψ is a homogeneous matrix concave function, its matrix hypograph $\mathcal{K}^n$ is a convex cone:

$$\mathcal{K}^n = \{(A, B, C) \in \operatorname{dom} \Psi \times \mathbf{S}^n : \Psi(A, B) \succeq C\}. \tag{6.33}$$

To obtain certificates of nonnegativity for entropy-like functions of the form (6.24) we are going to define a lifting map A from $\mathbb{R}^n_{++}$ to $\mathcal{K}^m$ for some m. Given $x = (x_1, \ldots, x_n)$, denote by $[x]_d$ the vector of all monomials up to degree d. If $i \in \{1, \ldots, n\}$ we can define the following map from $\mathbb{R}^n_{++}$ to $\mathcal{K}^m$ where $m = \dim \mathbb{R}[x_1, \ldots, x_n]_{\leq d-1}$.

$$A_i(x) = \left([x]_d [x]_d^T \ , \ x_i [x]_d [x]_d^T \ , \ \log(x_i) [x]_d [x]_d^T\right). \tag{6.34}$$

Note that even though the definition (6.32) assumes A and B to be invertible, the function Ψ can in fact be extended to the case where $\ker B \subseteq \ker A$, see e.g., [42]. In (6.34) the matrices $A = [x]_d [x]_d^T$ and $B = x_i [x]_d [x]_d^T$ commute in which case the value of $\Psi(A, B)$ can be shown to be $\log(x_i) [x]_d [x]_d^T$ which shows that $A_i(x) \in \mathcal{K}^m$ for all $x > 0$. Also note that all the entries of $A_i(x)$ are of the form $x^\alpha \log x_i$ and are thus of "entropy-like" form.

The dual cone to $\mathcal{K}^n$ is, by definition,

$$(\mathcal{K}^n)^* = \{(U,V,W) \in \mathbf{S}^n \times \mathbf{S}^n \times \mathbf{S}^n : \mathbf{Tr}[UA + VB + WC] \geq 0 \quad \forall (A,B,C) \in \mathcal{K}^n\}. \tag{6.35}$$

Unfortunately we do not have a closed form expression of this dual cone $\mathcal{K}^n$. In general gradients (and more generally Fenchel conjugates) of matrix functions are difficult to obtain due to the noncommuting nature of the arguments A and B.

Another important question is whether we can actually numerically solve convex optimization problems over (6.33) efficiently. Even though this cone is not natively supported by any solver we are aware of, one can obtain a sequence of tighter and tighter approximations of this cone using semidefinite programming, based on a limit formula similar to (6.31). In fact for the matrix perspective of log one can show that we have

$$\Psi(A,B) = \lim_{t \to 0} \frac{1}{t}(A \#_t B - A) \tag{6.36}$$

where $A \#_t B$ is the matrix perspective of the power function (also known as the t-weighted matrix geometric mean of A and B):

$$A \#_t B := A^{1/2} \left(A^{-1/2} B A^{-1/2}\right)^t A^{1/2}.$$

When t is rational, the function $(A,B) \mapsto A \#_t B$ admits a semidefinite programming formulation [35, 91]. By taking t small enough one can get a semidefinite approximation of the cone $\mathcal{K}^n$, via (6.36). These are the approximations we use for the numerical experiments presented in the next section.

6.6.4 Application: logarithmic Sobolev constants

In this section we use the certificates proposed here to compute lower bounds on the *logarithmic Sobolev constant* of any given finite Markov chain. These constants play an important role in bounding the mixing time of Markov chains [29].

We start by giving a brief description of the setting and notations needed to define these constants. We omit any discussion on the history and original motivation behind these constants and we refer instead to [29] for more information.

Setting Let $\mathcal{V}$ be a finite state-space and consider a Markov chain on $\mathcal{V}$ described by a transition matrix K, i.e., $K : \mathcal{V} \times \mathcal{V} \to \mathbb{R}_+$ is such that $K(u,v) \geq 0$ and $\sum_v K(u,v) = 1$ for all $u \in \mathcal{V}$. We assume that K is reversible with stationary distribution π, i.e., $\pi(u)K(u,v) = \pi(v)K(v,u)$ for all $u, v \in \mathcal{V} \times \mathcal{V}$. The Laplacian matrix associated to the Markov chain is

$$L = I - K. \tag{6.37}$$

and the differential equation describing the continuous-time Markov chain is:

$$\frac{d\mu}{dt} = -\mu L \tag{6.38}$$

where μ is a row vector representing a probability distribution on $\mathcal{V}$. Assuming that K is irreducible (i.e., that the undirected graph where edge uv has weight $\pi(u)K(u,v) = \pi(v)K(v,u)$ is connected) then for any initial distribution $\mu(0)$ on $\mathcal{V}$ we have $\mu(t) \to \pi$ when $t \to +\infty$. A main concern in the study of Markov chains is to study how fast this convergence happens. We now introduce two important quantities, namely the *spectral gap* and the *logarithmic Sobolev constant*, that can be used to bound the mixing time of this Markov chain.

Spectral gap Since (K, π) is reversible the Laplacian L defined in (6.37) is a self-adjoint map on $\mathbb{R}^{\mathcal{V}}$ for the inner product

$$\langle f, g \rangle_\pi = \sum_{u \in \mathcal{V}} \pi(u) f(u) g(u).$$

As such the eigenvalues of L are all real; furthermore they are all nonnegative since L is diagonally dominant. The smallest eigenvalue of L is zero since $L\,1 = 0$ where 1 is the all-ones vector. The *spectral gap* λ of the Markov chain K is defined as the second smallest eigenvalue of L. It has the following variational formulation:

$$\lambda = \min_{f \in \mathbb{R}^{\mathcal{V}}} \{ \langle f, Lf \rangle_\pi : \mathbb{E}_\pi[f] = 0,\ \mathbb{E}_\pi[f^2] = 1 \}. \tag{6.39}$$

The constraint $\mathbb{E}_\pi[f] = 0$ states that f is orthogonal to the all-ones vector 1, and $\mathbb{E}_\pi[f^2] = 1$ is equivalent to $\langle f, f \rangle_\pi = 1$. The spectral gap allows us to bound the rate at which μ converges to π. For example it is not difficult to show, using standard comparison theorems for dynamical systems, that for any initial distribution $\mu(0)$ on $\mathcal{V}$ the trajectory $\mu(t)$ satisfies

$$\|\mu(t) - \pi\|_{\mathrm{TV}}^2 \leq \frac{1}{4} \cdot \frac{1}{\pi_*} e^{-2\lambda t} \tag{6.40}$$

where $\|\mu - \pi\|_{\mathrm{TV}}$ is the total variation distance

$$\|\mu - \pi\|_{\mathrm{TV}} := \frac{1}{2} \sum_{u \in \mathcal{V}} |\mu(u) - \pi(u)|$$

and where $\pi_* = \min_{u \in \mathcal{V}} \pi(u)$ (note that the dependence on the initial distribution $\mu(0)$ does not appear in (6.40), and instead it was bounded above by the term $1/\pi_*$), see e.g., [29]. The spectral gap is probably the most popular way to bound the mixing time of a Markov chain. In some cases however one wants to resort to other techniques that yield more accurate bounds. One such technique is based on the logarithmic Sobolev constant.

Log-Sobolev constant Recall that if μ and π are two probability distributions on $\mathcal{V}$ then their KL-divergence is defined as:

$$D_{\mathrm{KL}}(\mu\|\pi) = \sum_{u\in\mathcal{V}} \mu(u)\log(\mu(u)/\pi(u)).$$

Definition 20 (Logarithmic Sobolev constant)**.** Let (K,π) be a finite reversible Markov chain and let $L = I - K$ be the associated Laplacian. The logarithmic Sobolev constant associated to (K,π) is defined as:

$$\alpha = \min_{f\in\mathbb{R}^{\mathcal{V}}} \left\{ \frac{\langle f, Lf\rangle_\pi}{D_{\mathrm{KL}}(f^2\pi\|\pi)} : \mathbb{E}_\pi[f^2] = 1 \right\}. \tag{6.41}$$

Equivalently, it is the largest constant c such that $\langle f, Lf\rangle_\pi \geq cD_{\mathrm{KL}}(f^2\pi\|\pi)$ holds for all $f\in\mathbb{R}^{\mathcal{V}}$ satisfying $\mathbb{E}_\pi[f^2] = 1$.

Note that in (6.41) the expression $f^2\pi$ corresponds to the probability distribution $\mu(u) = f^2(u)\pi(u)$ (this is a valid probability distribution since $\mathbb{E}_\pi[f^2] = 1$). The importance of the constant α comes from the fact that it gives a bound on the mixing time of the Markov chain (K,π). Indeed it is shown in [29] that for any initial distribution $\mu(0)$ on $\mathcal{V}$ the solution $\mu(t)$ of (6.38) satisfies:

$$\|\mu(t) - \pi\|_{\mathrm{TV}}^2 \leq \frac{1}{2}\log\left(\frac{1}{\pi_*}\right) e^{-2\alpha t}. \tag{6.42}$$

Note that if $\alpha \approx \lambda$, the estimate (6.42) can be much better than (6.40), especially when the state space has exponential size. Indeed if for example $\mathcal{V} = \{0,1\}^n$ and π_* is the uniform distribution then the coefficient in (6.40) is $1/\pi_* = 2^n$ whereas in (6.42) it is $\log(1/\pi_*) = n$.

Lower bound on α using convex optimization Unfortunately the problem of computing the logarithmic Sobolev constant of a given Markov chain is very difficult [92, page 336] (we are not aware though of any formal complexity result concerning the computation of α). The objective of this section is to show how the certificates presented earlier can be used to compute a lower bound on α. We also present numerical experiments indicating that the bounds we get are very close to the correct values of α (when it is known).

First, observe that the problem of computing α can be formulated as the problem of certifying nonnegativity of an entropy-like function. Indeed if we expand the definition of $\langle f, Lf\rangle_\pi$ and of $D_{\mathrm{KL}}(f^2\pi\|\pi)$ in (6.41), and use the fact that f can be assumed positive (since $\langle |f|, L|f|\rangle_\pi \leq \langle f, Lf\rangle_\pi$) we see that α is the largest constant c such that the following expression

$$\frac{1}{2}\sum_{u,v\in\mathcal{V}} \pi(u)K(u,v)(f(u) - f(v))^2 - 2c\sum_{u\in\mathcal{V}} f(u)^2\pi(u)\log(f(u)) \tag{6.43}$$

is nonnegative on $\{f\in\mathbb{R}^{\mathcal{V}}_{++} : \mathbb{E}_\pi[f^2] = 1\}$. Note that the expression (6.43) (where

the variables are the $f(u)$, $u \in \mathcal{V}$) is an "entropy-like" function of the form (6.24). Let $A_u : \mathbb{R}^n \to \mathcal{K}^{n+1}$ for $u \in \mathcal{V}$ be the lifting map defined in (6.34) with $d = 1$:

$$A_u(x) = \Big([x]_1 [x]_1^T \ , \ x_u [x]_1 [x]_1^T \ , \ \log(x_u) [x]_1 [x]_1^T \Big).$$

We can use this lifting map, together with the sum-of-squares proof system, to define a lower bound α_{cvx} on α that can be computed using convex optimization. This is the object of the next definition.

Definition 21. Given a reversible Markov chain (K, π) on state space $\{1, \ldots, n\}$ we define α_{cvx} to be the largest constant c such that there exist $B_u \in (\mathcal{K}^{n+1})^*$ for $u \in \{1, \ldots, n\}$, $sos \in \text{SOS}_{n,4}$ (a sum-of-squares polynomial of degree at most 4), and $p \in \mathbb{R}[x_1, \ldots, x_n]_{\leq 2}$ such that the following identity holds for all $x \in \mathbb{R}^n_{++}$:

$$\begin{aligned} &\frac{1}{2} \sum_{1 \leq u,v \leq n} \pi(u) K(u,v) (x_u - x_v)^2 - 2c \sum_{u \in \mathcal{V}} x_u^2 \pi(u) \log(x_u) \\ &\qquad = \sum_{1 \leq u \leq n} \langle A_u(x), B_u \rangle + sos(x) + p(x) \left(\sum_{1 \leq u \leq n} \pi(u) x_u^2 - 1 \right). \end{aligned} \tag{6.44}$$

Note that we used the notation x_u instead of $f(u)$ to denote the indeterminates in (6.44), to make it more consistent with the previous sections. Since the right-hand side of (6.44) is nonnegative whenever $\sum_{1 \leq u \leq n} \pi(u) x_u^2 = 1$ we see that the constant α_{cvx} satisfies $\alpha_{\text{cvx}} \leq \alpha$ where α is the log-Sobolev constant of (K, π). Furthermore α_{cvx} can be computed as the solution of a convex optimization problem involving the positive semidefinite cone (for the *sos* term) and the cone $\mathcal{K}^{n+1}$ that we introduced in the previous section.

Numerical examples We tested our approach on some examples of transition matrices K for which the log-Sobolev constant is known, and some where it is unknown. More specifically we looked at the following three examples of Markov chain:

- **Two-point space**: This is a Markov chain on a two-point space $\mathcal{V} = \{-1, 1\}$ with the transition matrix
$$K = \begin{bmatrix} \theta & 1-\theta \\ \theta & 1-\theta \end{bmatrix}$$
where $\theta \in (0, 1/2]$. This chain is reversible with respect to the stationary distribution $\pi = (\theta, 1-\theta)$. The spectral gap of this chain is equal to 1, and the log-Sobolev constant, computed in [29, Theorem A.2] is equal to $(1-\theta)/\log((1-\theta)/\theta)$.

- **Complete graph**: This is a Markov chain on the complete graph where each node has probability $1/(n-1)$ of transitioning to one of its neighbors. The transition matrix is thus $K = \frac{1}{n-1}(1\,1^T - I_n)$ and the stationary distribution is uniform $\pi = 1\,/n$. The log-Sobolev constant was shown in [29, Corollary A.5] to be equal to $\frac{n-2}{(n-1)\log(n-1)}$.

- **Cycle graph**: Finally the last example we consider is the cycle graph on n nodes where each node has probability $1/2$ of transitioning to one of its two neighbors. Thus $K(u, u \pm 1) = \frac{1}{2}$ for $u \in \mathbb{Z}_n$ and 0 otherwise. The stationary distribution is the uniform distribution. The log-Sobolev constant was shown in [26] to be equal to $\frac{1}{2}(1 - \cos(2\pi/n))$ when n is even. However the value for n odd is unknown.

Table 6.1 shows the numerical results obtained on the different chains discussed above. To compute α_{cvx} we used the approximation given by the limit formula (6.36), with the semidefinite programming formulation of the matrix geometric mean given in [35] with $t = 2^{-6}$. The computations were done on Matlab using CVX [54, 53] and the solver Sedumi.

Markov chain		α_{cvx} (approx. (6.36), $t = 2^{-6}$)	Exact α (if known)	$\lambda/2$
Two-point space [29, Thm. A.2] $K = \left[\begin{smallmatrix} \theta & 1-\theta \\ \theta & 1-\theta \end{smallmatrix}\right]$ $\pi = (\theta, 1-\theta)$ $\alpha = \frac{(1-\theta)}{\log((1-\theta)/\theta)}$	$\theta = 0.1$	0.3602	0.3641	0.5000
	$\theta = 0.2$	0.4289	0.4328	0.5000
	$\theta = 0.3$	0.4682	0.4721	0.5000
	$\theta = 0.4$	0.4894	0.4933	0.5000
	$\theta = 0.5$	0.4969	0.5000	0.5000
Complete graph [29, Cor. A.5] $K = \frac{1}{n-1}(1\,1^T - I_n)$ $\pi = 1/n$ $\alpha = \frac{n-2}{(n-1)\log(n-1)}$	$n = 3$	0.7155	0.7213	0.7500
	$n = 4$	0.6017	0.6068	0.6667
	$n = 5$	0.5362	0.5410	0.6250
	$n = 6$	0.4924	0.4971	0.6000
	$n = 7$	0.4606	0.4651	0.5833
	$n = 8$	0.4360	0.4405	0.5714
Cycle graph [26] For n even $\alpha = \frac{1}{2}(1 - \cos\left(\frac{2\pi}{n}\right))$	$n = 4$	0.4970	0.5000	0.5000
	$n = 5$	0.3446	?	0.3455
	$n = 6$	0.2491	0.2500	0.2500
	$n = 7$	0.1878	?	0.1883
	$n = 8$	0.1465*	0.1464	0.1464

Table 6.1: Numerical lower bounds on the log-Sobolev constants for some finite Markov chains. The last column gives the value of $\lambda/2$ (λ is the spectral gap) which is always an upper bound on the log-Sobolev constant [92, Lemma 2.2.2].
*The quantity α_{cvx} is always a lower bound to α however numerical errors can cause small violations of the inequality $\alpha_{\text{cvx}} \leq \alpha$.

Comments and open questions We see from Table 6.1 that α_{cvx} gets very close to the true value of α for the different chains that were considered. There are several open questions concerning α_{cvx} that it would be interesting to explore further:

- *Tensorization:* One important property of the log-Sobolev constant α is *tensorization*. Assume (K_1, π_1) and (K_2, π_2) are two reversible Markov chains on state spaces $\mathcal{V}_1$ and $\mathcal{V}_2$. Consider the *product* Markov chain K on state space $\mathcal{V} = \mathcal{V}_1 \times \mathcal{V}_2$ defined as follows: from state $(u_1, u_2) \in \mathcal{V}_1 \times \mathcal{V}_2$ we choose a coordinate $i = 1$ or 2 with probability $1/2$ and we update the i'th component using transition matrix K_i. In matrix terms, the transition matrix K of the product chain is given by:

$$K = \frac{1}{2}(K_1 \otimes I + I \otimes K_2)$$

  where $\otimes$ denotes Kronecker product. If we let $\pi(u_1, u_2) = \pi_1(u_1)\pi_2(u_2)$ then it is not hard to see that (K, π) is reversible. The remarkable property about the log-Sobolev constant of (K, π) is that it satisfies:

$$\alpha = \frac{1}{2}\min(\alpha_1, \alpha_2)$$

  where α_1, α_2 are respectively the log-Sobolev constants of (K_1, π_1) and (K_2, π_2). This property allows us for example to get the log-Sobolev constant of the random walk on the hypercube $\{0, 1\}^n$ where at each step we flip one coordinate chosen at random, from the log-Sobolev constant of the two-point space. As this is a crucial property of α an important question is to know whether the relaxation α_{cvx} also satisfies such a tensorization property.

- *Modified log-Sobolev constants:* There are other constants similar to α that have been proposed to quantify the mixing time of Markov chains and are sometimes called modified log-Sobolev constants [11]. These constants can give sharper estimates about the convergence of the Markov chain to the stationary distribution. One interesting question is to know whether the approach proposed here can be applied also to compute modified logarithmic Sobolev constants.

- *Fastest Markov chain in terms of log-Sobolev constant:* In [14] Boyd, Diaconis and Xiao studied the following problem: given an unweighted undirected graph G find a symmetric transition probability matrix such that the resulting Markov chain (with the uniform stationary distribution) has the "fastest" mixing, in terms of the spectral gap. The resulting optimization problem can be shown to be a semidefinite program [14]. A natural question is to study the same problem where the notion of "fastest" is measured according to the log-Sobolev constant, rather than the spectral gap (in fact the authors of [14] pose this question at the end of their paper). It would be interesting to know whether the fastest Markov chains in both cases can be very different from each other.

6.7 Summary of chapter

- We consider the general problem of certifying global nonnegativity of a function f defined on some set X.
- A conic certificate of nonnegativity on X is given by the choice of a convex cone K and a lifting map $A : X \to K$. Given $f : X \to \mathbb{R}$ one can try to certify that f is nonnegative by finding $B \in K^*$ such that

$$f(x) = \langle A(x), B \rangle \quad \forall x \in X.$$

- Many existing certificates of nonnegativity fall into this framework: LP certificates (Farkas, Krivine, Handelman, Sherali-Adams), SOS certificates, certificates based on geometric programming [46], certificates for signomial functions [24].
- We use this framework to develop a new way to certify nonnegativity of entropy-like functions, i.e., functions of the form $p_0(x) + \sum_{i=1}^{n} p_i(x) \log(x_i)$. Our certificates exploit the matrix concavity of the logarithm function. As an application we show how it can be used to numerically estimate the logarithm Sobolev constant of finite Markov chains.

Bibliography

[1] Jim Agler, William Helton, Scott McCullough, and Leiba Rodman. Positive semidefinite matrices with a given sparsity pattern. *Linear algebra and its applications*, 107:101–149, 1988. 109

[2] Boaz Barak and David Steurer. Sum-of-squares proofs and the quest toward optimal algorithms. In *Proceedings of International Congress of Mathematicians (ICM)*, 2014. 20, 27, 28

[3] Alexander Barvinok. *A course in convexity*, volume 54. American Mathematical Society, 2002. 103

[4] Alexander Barvinok and Grigoriy Blekherman. Convex geometry of orbits. *Combinatorial and Computational Geometry. MSRI Publications*, 52:51–77, 2005. 63

[5] Alexander Barvinok and Anatolii Moiseevich Vershik. Convex hulls of orbits of representations of finite groups and combinatorial optimization. *Functional Analysis and Its Applications*, 22(3):224–225, 1988. 63

[6] Aharon Ben-Tal and Arkadi Nemirovski. On polyhedral approximations of the second-order cone. *Mathematics of Operations Research*, 26(2):193–205, 2001. 14

[7] Abraham Berman and Naomi Shaked-Monderer. *Completely positive matrices*. World Scientific Pub Co Inc, 2003. 46

[8] Rajendra Bhatia. *Positive definite matrices*. Princeton University Press, 2009. 146

[9] Grigoriy Blekherman, João Gouveia, and James Pfeiffer. Sums of squares on the hypercube. *arXiv preprint arXiv:1402.4199*, 2014. 103

[10] Grigoriy Blekherman, Pablo A. Parrilo, and Rekha R. Thomas. *Semidefinite optimization and convex algebraic geometry*. SIAM, 2013. 9, 28, 59, 136

[11] Sergey G. Bobkov and Prasad Tetali. Modified logarithmic Sobolev inequalities in discrete settings. *Journal of Theoretical Probability*, 19(2):289–336, 2006. 153

[12] Yuri Bogomolov, Samuel Fiorini, Aleksandr Maksimenko, and Kanstantsin Pashkovich. Small extended formulations for cyclic polytopes. *Discrete & Computational Geometry*, 53(4):809–816, 2015. 105, 125

[13] Jonathan Borwein and Henry Wolkowicz. Regularizing the abstract convex program. *Journal of Mathematical Analysis and Applications*, 83(2):495–530, 1981. 40

[14] Stephen Boyd, Persi Diaconis, and Lin Xiao. Fastest mixing Markov chain on a graph. *SIAM Review*, 46(4):667–689, 2004. 153

[15] Gábor Braun, Samuel Fiorini, Sebastian Pokutta, and David Steurer. Approximation limits of linear programs (beyond hierarchies). In *IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 480–489, 2012. 9

[16] Gábor Braun, Rahul Jain, Troy Lee, and Sebastian Pokutta. Information-theoretic approximations of the nonnegative rank. *To Appear in Computational complexity*, pages 1–51, 2016. 35

[17] Gábor Braun and Sebastian Pokutta. Common information and unique disjointness. In *IEEE 54th Annual Symposium on Foundations of Computer Science*, 2013. 35

[18] Mark Braverman and Ankur Moitra. An information complexity approach to extended formulations. In *Proceedings of the 45th Annual ACM symposium on Symposium on Theory of Computing*, pages 161–170. ACM, 2013. 9, 35

[19] Eric Carlen. Trace inequalities and quantum entropy: an introductory course. *Entropy and the quantum*, 529:73–140, 2010. 146

[20] Robert D. Carr and Goran Konjevod. Polyhedral combinatorics. In *Tutorials on emerging methodologies and applications in Operations Research, chapter 2, edited by Harvey Greenberg*. Springer, 2004. 70

[21] Siu On Chan, James R. Lee, Prasad Raghavendra, and David Steurer. Approximate constraint satisfaction requires large LP relaxations. In *IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 350–359, 2013. 9, 57, 79

[22] Venkat Chandrasekaran, Benjamin Recht, Pablo A. Parrilo, and Alan S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012. 36

[23] Venkat Chandrasekaran and Parikshit Shah. Relative entropy optimization and its applications. *Mathematical Programming*, 2015. 140

[24] Venkat Chandrasekaran and Parikshit Shah. Relative entropy relaxations for signomial optimization. *SIAM Journal on Optimization*, 26(2):1147–1173, 2016. 132, 139, 140, 154

[25] Robert Chares. *Cones and interior-point algorithms for structured convex optimization involving powers and exponentials*. PhD thesis, Université Catholique de Louvain, Louvain-la-Neuve, 2008. 137, 138

[26] Guan-Yu Chen and Yuan-Chung Sheu. On the log-sobolev constant for the simple random walk on the n-cycle: the even cases. *Journal of Functional Analysis*, 202(2):473–485, 2003. 152

[27] Michele Conforti, Gérard Cornuéjols, and Giacomo Zambelli. Extended formulations in combinatorial optimization. *4OR*, 8(1):1–48, 2010. 14

[28] Etienne de Klerk and Frank Vallentin. On the Turing model complexity of interior point methods for semidefinite programming. *arXiv preprint arXiv:1507.03549*, 2015. 29

[29] Persi Diaconis and Laurent Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *The Annals of Applied Probability*, 6(3):695–750, 1996. 148, 149, 150, 151, 152

[30] Mathias Drton, Bernd Sturmfels, and Seth Sullivant. *Lectures on algebraic statistics*. Springer, 2009. 32

[31] Bogdan Dumitrescu. *Positive trigonometric polynomials and signal processing applications*. Springer, 2007. 103

[32] Hamza Fawzi, João Gouveia, Pablo A. Parrilo, Richard Z. Robinson, and Rekha R. Thomas. Positive semidefinite rank. *Mathematical Programming*, 153(1):133–177, 2015. 23

[33] Hamza Fawzi and Pablo A. Parrilo. Lower bounds on nonnegative rank via nonnegative nuclear norms. *Mathematical Programming*, 153(1):41–66, 2015. 38, 45

[34] Hamza Fawzi and Pablo A. Parrilo. Self-scaled bounds for atomic cone ranks: applications to nonnegative rank and cp-rank. *Mathematical Programming*, 158(1):417–465, 2016. 10, 31, 47

[35] Hamza Fawzi and James Saunderson. Lieb's concavity theorem, matrix geometric means, and semidefinite optimization. *arXiv preprint arXiv:1512.03401*, 2015. 148, 152

[36] Hamza Fawzi, James Saunderson, and Pablo A. Parrilo. Equivariant semidefinite lifts of regular polygons. *To Appear in Mathematics of Operations Research*, 2014. 10, 56, 86, 97

[37] Hamza Fawzi, James Saunderson, and Pablo A. Parrilo. Equivariant semidefinite lifts and sum-of-squares hierarchies. *SIAM Journal on Optimization*, 25(4):2212–2243, 2015. 10, 56

[38] Hamza Fawzi, James Saunderson, and Pablo A. Parrilo. Sparse sums of squares on finite abelian groups and improved semidefinite lifts. *To Appear in Mathematical Programming*, 2015. 9, 10, 97

[39] Carla Fidalgo and Alexander Kovacec. Positive semidefinite diagonal minus tail forms are sums of squares. *Mathematische Zeitschrift*, 269(3-4):629–645, 2011. 136

[40] Samuel Fiorini, Volker Kaibel, Kanstantsin Pashkovich, and Dirk Oliver Theis. Combinatorial bounds on nonnegative rank and extended formulations. *Discrete Mathematics*, 313(1):67 – 83, 2013. 33, 34, 105, 124

[41] Samuel Fiorini, Serge Massar, Sebastian Pokutta, Hans R. Tiwary, and Ronald de Wolf. Linear vs. semidefinite extended formulations: exponential separation and strong lower bounds. In *Proceedings of the 44th Symposium on Theory of Computing*, pages 95–106. ACM, 2012. 9, 29, 34

[42] J.I. Fujii. Operator means and the relative operator entropy. In *Operator Theory and Complex Analysis*, pages 161–172. Springer, 1992. 147

[43] J.I. Fujii and Eizaburo Kamei. Relative operator entropy in noncommutative information theory. *Math. Japon*, 34:341–348, 1989. 147

[44] Shigeru Furuichi, Kenjiro Yanagi, and Ken Kuriyama. A note on operator inequalities of Tsallis relative operator entropy. *Linear Algebra and its Applications*, 407:19–31, 2005. 147

[45] David Gale. Neighborly and cyclic polytopes. In *Proceedings of the Seventh Symposium in Pure Mathematics of the American Mathematical Society*, volume 7, pages 225–232, 1963. 105, 124

[46] Mehdi Ghasemi and Murray Marshall. Lower bounds for polynomials using geometric programming. *SIAM Journal on Optimization*, 22(2):460–473, 2012. 132, 136, 138, 139, 154

[47] Michel X. Goemans. Smallest compact formulation for the permutahedron. *Mathematical Programming*, 153(1):5–11, 2015. 9, 14, 57

[48] João Gouveia, Monique Laurent, Pablo A. Parrilo, and Rekha R. Thomas. A new semidefinite programming hierarchy for cycles in binary matroids and cuts in graphs. *Mathematical programming*, 133(1-2):203–225, 2012. 73, 74

[49] João Gouveia, Pablo A. Parrilo, and Rekha R. Thomas. Theta bodies for polynomial ideals. *SIAM Journal on Optimization*, 20(4):2097–2118, 2010. 28

[50] João Gouveia, Pablo A. Parrilo, and Rekha R. Thomas. Lifts of convex sets and cone factorizations. *Mathematics of Operations Research*, 38(2):248–264, 2013. 10, 12, 21, 22, 23, 57, 63, 64, 132, 133

[51] João Gouveia, Richard Z. Robinson, and Rekha R. Thomas. Polytopes of minimum positive semidefinite rank. *Discrete & Computational Geometry*, 50(3):679–699, 2013. 69

[52] João Gouveia and Rekha Thomas. Convex hulls of algebraic sets. In *Handbook on Semidefinite, Conic and Polynomial Optimization*, pages 113–138. Springer, 2012. 9

[53] Michael Grant and Stephen Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. `http://stanford.edu/~boyd/graph_dcp.html`. 152

[54] Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. `http://cvxr.com/cvx`, March 2014. 152

[55] Andreas Griewank and Philippe L. Toint. On the existence of convex decompositions of partially separable functions. *Mathematical Programming*, 28(1):25–49, 1984. 106, 109

[56] Robert Grone, Charles R. Johnson, Eduardo M. Sá, and Henry Wolkowicz. Positive definite completions of partial hermitian matrices. *Linear algebra and its applications*, 58:109–124, 1984. 53, 106, 109, 110

[57] Martin Grötschel, László Lovász, and Alexander Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981. 7, 29

[58] David Handelman. Representing polynomials by positive linear functions on compact convex polyhedra. *Pacific Journal of Mathematics*, 132(1):35–62, 1988. 28, 135

[59] J. William Helton and Jiawang Nie. Sufficient and necessary conditions for semidefinite representability of convex hulls and sets. *SIAM Journal on Optimization*, 20(2):759–791, 2009. 9

[60] R. G. Jeroslow. On defining sets of vertices of the hypercube by linear inequalities. *Discrete Mathematics*, 11(2):119–124, 1975. 70

[61] V. Kaibel. Extended formulations in combinatorial optimization. *Arxiv preprint arXiv:1104.1023, Appeared in Optima 85 newsletter of the Mathematical Optimization Society*, 2011. 14

[62] Volker Kaibel and Kanstantsin Pashkovich. Constructing extended formulations from reflection relations. In *Integer Programming and Combinatoral Optimization*, pages 287–300. Springer, 2011. 9, 70, 125

[63] Volker Kaibel, Kanstantsin Pashkovich, and Dirk Theis. Symmetry matters for the sizes of extended formulations. *Integer programming and combinatorial optimization*, pages 135–148, 2010. 57

[64] Volker Kaibel, Kanstantsin Pashkovich, and Dirk Oliver Theis. Symmetry matters for sizes of extended formulations. *SIAM Journal on Discrete Mathematics*, 26(3):1361–1382, 2012. 9

[65] Mauricio Karchmer, Eyal Kushilevitz, and Noam Nisan. Fractional covers and communication complexity. In *Proceedings of the Seventh Annual Structure in Complexity Theory Conference.*, pages 262–274. IEEE, 1992. 42

[66] Jean-Louis Krivine. Quelques propriétés des préordres dans les anneaux commutatifs unitaires. *Comptes Rendus de l'Académie des Sciences de Paris*, 258:3417–3148, 1964. 28, 135

[67] Kaie Kubjas, Elina Robeva, Bernd Sturmfels, et al. Fixed points EM algorithm and nonnegative rank boundaries. *The Annals of Statistics*, 43(1):422–461, 2015. 32

[68] Jean-Bernard Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001. 111

[69] Jean-Bernard Lasserre. Convex sets with semidefinite representation. *Mathematical Programming*, 120(2):457–477, 2009. 9, 28

[70] Jean-Bernard Lasserre. *An Introduction to Polynomial and Semi-Algebraic Optimization*. Cambridge University Press, 2015. 136

[71] Monique Laurent. A comparison of the sherali-adams, Lovász-Schrijver, and Lasserre relaxations for 0–1 programming. *Mathematics of Operations Research*, 28(3):470–496, 2003. 28

[72] Monique Laurent. Lower bound for the number of iterations in semidefinite hierarchies for the cut polytope. *Mathematics of Operations Research*, 28(4):871–883, 2003. 10, 76, 77, 79, 103, 104, 115, 119, 125

[73] Monique Laurent. Semidefinite relaxations for max-cut. In *The Sharpest Cut*, chapter 16, pages 257–290. SIAM, 2004. 77

[74] Monique Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging applications of algebraic geometry*, pages 157–270. Springer, 2009. 28

[75] Monique Laurent and Zhao Sun. Handelman's hierarchy for the maximum stable set problem. *Journal of Global Optimization*, 60(3):393–423, 2014. 28, 135

[76] James R. Lee, Prasad Raghavendra, and David Steurer. Lower bounds on the size of semidefinite programming relaxations. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC)*, pages 567–576. ACM, 2015. 9, 29

[77] James R. Lee, Prasad Raghavendra, David Steurer, and Ning Tan. On the power of symmetric LP and SDP relaxations. http://www.cs.cornell.edu/~dsteurer/papers/symsdp/, 2014. 59

[78] Troy Lee and Adi Shraibman. Lower bounds in communication complexity. *Foundations and Trends®in Theoretical Computer Science*, 3(4):263–399, 2009. 32

[79] Cong Han Lim and Stephen Wright. Beyond the birkhoff polytope: Convex relaxations for vector permutation problems. In *Advances in Neural Information Processing Systems*, pages 2168–2176, 2014. 13

[80] László Lovász. On the ratio of optimal integral and fractional covers. *Discrete mathematics*, 13(4):383–390, 1975. 42

[81] László Lovász. Communication complexity: A survey. In B. Korte, H.J. Promel, and R. L. Graham, editors, *Paths, Flows, and VLSI-Layout*, pages 235–265. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1990. 32

[82] C. Lund and M. Yannakakis. On the hardness of approximating minimization problems. *Journal of the ACM*, 41(5):960–981, 1994. 42

[83] Yurii Nesterov. Squared functional systems and optimization problems. In *High performance optimization*, pages 405–440. Springer, 2000. 111

[84] Yurii Nesterov and Michael J. Todd. Self-scaled barriers and interior-point methods for convex programming. *Mathematics of Operations Research*, 22(1):1–42, 1997. 38

[85] Pablo A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, 2000. 111

[86] Kanstantsin Pashkovich. Tight lower bounds on the sizes of symmetric extensions of permutahedra and similar results. *arXiv preprint arXiv:0912.3446*, 2009. 57

[87] Frank Permenter and Pablo Parrilo. Partial facial reduction: simplified, equivalent SDPs via approximations of the PSD cone. *arXiv preprint arXiv:1408.4685*, 2014. 40

[88] R Tyrell Rockafellar. *Convex analysis*, volume 28. Princeton University Press, 1997. 37

[89] Thomas Rothvoß. The matching polytope has exponential extension complexity. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 263–272. ACM, 2014. 9, 29, 34

[90] Walter Rudin. *Fourier Analysis on Groups*. John Wiley & Sons, Inc., 1990. 107

[91] Guillaume Sagnol. On the semidefinite representation of real functions applied to symmetric matrices. *Linear Algebra and its Applications*, 439(10):2829–2843, 2013. 148

[92] Laurent Saloff-Coste. Lectures on finite Markov chains. In *Lectures on probability theory and statistics*, pages 301–413. Springer, 1997. 150, 152

[93] Raman Sanyal, Frank Sottile, and Bernd Sturmfels. Orbitopes. *Mathematika*, 57(02):275–314, 2011. 63

[94] James Saunderson, Pablo A. Parrilo, and Alan S. Willsky. Semidefinite descriptions of the convex hull of rotation matrices. *arXiv preprint arXiv:1403.4914*, 2014. 9

[95] Jean-Pierre Serre. *Linear representations of finite groups*. Springer New York, 1977. 60, 81

[96] Hanif D. Sherali and Warren P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3(3):411–430, 1990. 27, 135

[97] Audrey Terras. *Fourier analysis on finite groups and applications*. Cambridge University Press, 1999. 107

[98] Joel A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015. 147

[99] Levent Tunçel. Potential reduction and primal-dual methods. In *Handbook of semidefinite programming*, pages 235–265. Springer, 2000. 57

[100] Aaron Wyner. The common information of two dependent random variables. *IEEE Transactions on Information Theory*, 21(2):163–179, 1975. 35

[101] Mihalis Yannakakis. Expressing combinatorial optimization problems by linear programs. *Journal of Computer and System Sciences*, 43(3):441–466, 1991. 9, 10, 12, 15, 17, 57, 70

[102] Günter M. Ziegler. *Lectures on polytopes*, volume 152. Springer, 1995. 103, 124