Report from Dagstuhl Seminar 15082

Limitations of Convex Programming: Lower Bounds on Extended Formulations and Factorization Ranks

Edited by

Hartmut Klauck[1], Troy Lee[2], Dirk Oliver Theis[3], and Rekha R. Thomas[4]

1 **Nanyang Technological University, SG,** cqthk@nus.edu.sg
2 **Nanyang Technological University, SG,** troyjlee@gmail.com
3 **University of Tartu, EE,** dirk.oliver.theis@ut.ee
4 **University of Washington, US,** rrthomas@uw.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 15082 "Limitations of convex programming: lower bounds on extended formulations and factorization ranks" held in February 2015. Summaries of a selection of talks are given in addition to a list of open problems raised during the seminar.

Seminar February 15–20, 2015 – http://www.dagstuhl.de/15082
1998 ACM Subject Classification G.1.6 Convex Optimization, F.2 Analysis of Algorithms and Problem Complexity
Keywords and phrases Convex optimization, extended formulations, cone rank, positive semidefinite rank, nonnegative rank, quantum communication complexity, real algebraic geometry
Digital Object Identifier 10.4230/DagRep.5.2.109

1 Executive Summary

Hartmut Klauck
Troy Lee
Dirk Oliver Theis
Rekha R. Thomas

The topic of this seminar was the rapidly developing notion of *cone rank* of a matrix/polytope that is an important invariant controlling several properties of the matrix/polytope with connections to optimization, communication complexity and theoretical computer science. This meeting was a follow-up to the 2013 Dagstuhl seminar *13082: "Communication Complexity, Linear Optimization, and lower bounds for the nonnegative rank of matrices"* organized by Leroy Beasley, Hartmut Klauck, Troy Lee and Dirk Oliver Theis.

The cone rank of a nonnegative matrix is an ordered notion of matrix rank with emerging applications in several fields. A well-known example is the nonnegative rank of a nonnegative matrix that appears in areas ranging from communication complexity, to statistics, to combinatorial optimization, and algebraic complexity theory. A related notion arising as an invariant in representations of convex sets as projections of affine slices of the positive semidefinite (psd) cone is the positive semidefinite (psd) rank. The psd rank is a very new quantity and is still relatively poorly understood.

Limitations of Convex Programming: Lower Bounds on Extended Formulations and Factorization Ranks,
Dagstuhl Reports, Vol. 5, Issue 2, pp. 109–127
Editors: Hartmut Klauck, Troy Lee, Dirk Oliver Theis, and Rekha R. Thomas
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The purpose of this seminar was to bring together researchers from optimization, computer science, real/convex/tropical algebraic geometry, and matrix theory, to discuss relevant techniques from each area that can contribute to the development of both the theory and computation of cone ranks and cone factorizations of nonnegative matrices, as well as their many emerging applications.

In optimization and computer science, a common approach to finding approximate solutions to NP-hard problems is to look at tractable convex relaxations of the problem as either a linear or semidefinite program. An optimal solution to such a relaxation gives a bound on the objective value of the original problem. While much previous work has focused on specific relaxations of a problem, or a family of relaxations coming from a hierarchy, cone ranks allow the study of the best possible linear, semidefinite, or other convex formulations of a NP-hard problem independent of specific construction methods. These formulations all write the underlying feasible set as the projection of an affine slice of a closed convex cone and is commonly refereed to as an extended formulation of the underlying feasible region. The nonnegative rank of a polytope is the smallest size of a linear extended formulation of the polytope while psd rank of the polytope is the size of the smallest possible semidefinite extended formulation of the polytope. Linear extended formulations are the best understood so far and an exciting development in this area is the recent breakthrough by Rothvoß showing that the matching polytope does not admit a polynomial sized linear extended formulation, settling a notorious open problem in combinatorial optimization. Very recently, there has also been exciting new developments in the area of psd rank such as the result of Lee, Raghavendra, and Steurer that shows that the psd rank of certain polytopes such as the traveling salesman polytope of a graph with n vertices must be exponential in n.

In the field of communication complexity, nonnegative and psd ranks are exactly characterized by a model of randomized and quantum communication complexity, respectively. This connection has allowed tools from communication and information theory to help create lower bounds for these ranks.

A central question in the field of real algebraic geometry is the semidefinite representability of convex sets. While polytopes only project to polytopes, affine slices of psd cones have much greater expressive power as their projections are convex sets, which allows the definition of psd rank for semi algebraic convex sets. Psd rank has inherent semi algebraic structure and its study crosses over into real and convex algebraic geometry, algebraic complexity, and semidefinite programming.

2 Table of Contents

3 Organization of the Seminar

One of the key goals of this seminar was to bring experts from all the different communities that intersect at the notion of cone ranks to inform the participants on the key tools and questions from each area that pertain to this topic. To achieve our goals of making the seminar part tutorial, we invited nine prominent researchers from a variety of areas to give expository lectures. These lectures were given by the following researchers listed along with the general topic of their talks:

1. Alexander Barvinok (University of Michigan): *Approximations of convex bodies by polytopes*
2. Greg Blekherman (Georgia Institute of Technology): *Sums of squares on the hypercube*
3. Hamza Fawzi (MIT): *Equivariant lifts of polytopes*
4. João Gouveia (University of Coimbra): *Survey on psd rank*
5. Volker Kaibel (University of Magdeburg): *Survey on nonnegative rank*
6. James Lee (University of Washington): *Lower bounds for SDP relaxations I*
7. Markus Schweighofer (University of Konstanz): *Real algebraic geometry*
8. David Steurer (Cornell University): *Lower bounds for SDP relaxations II*
9. Ronald de Wolf (CWI, Amsterdam): *Quantum communication complexity and psd rank*

Roughly two of these speakers were scheduled each day. In addition, there were short scientific talks by almost all the remaining participants as well as an open problems session on Wednesday evening.

4 Summaries of selected talks

In this section, we highlight some selected talks from the seminar. The first five are tutorial talks while the last is a sample research talk.

4.1 João Gouveia

João Gouveia gave the opening talk of the seminar which was a *survey on positive semidefinite rank*. The purpose of this talk was to get the audience primed for this topic which was a main feature at the seminar. The talk was based on the recent survey (with this title) that was written by Hamza Fawzi, João Gouveia, Pablo Parrilo, Richard Robinson and Rekha Thomas that is available on the arXiv.

The talk began with the definition of psd rank of a nonnegative matrix, basic properties and relationships to other common notions of rank. There was an instructive running example using circulant matrices that illustrated the main ranks that were introduced. Then he moved on to more sophisticated properties of psd rank such as the guarantee of factorizations of controlled norm. This result in particular played an important role in the recent work of Briët, Dadush and Pokutta who proved that 0/1 polytopes in a fixed dimension cannot all have small psd rank. Other features were also examined such as the space of factorizations, its connectivity properties, and symmetric factorizations. Several open questions were stated during the course of the talk which set the stage for the rest of the week. In particular, there was some discussion of the many open problems on the computational complexity of psd rank.

4.2 Greg Blekherman

Greg Blekherman gave another of our expository lectures on the topic *Symmetry Reduction for Sums of Squares on the Hypercube*.

The polytopes whose psd rank has attracted the most attention come from combinatorial optimization. In this setting, the polytope under consideration is the convex hull of a collection of vectors from $\{0,1\}^n$, where each vector denotes a subset of some ground set, and the problem is to minimize a polynomial function over it. This model covers many NP-hard problems and hence their psd rank is believed to be high. Yet, we have so far been able to only prove this for specific classes of polytopes such as cut polytopes (c.f. talks of James Lee and David Steurer).

The max cut problem which is NP-hard can be formulated as optimizing a quadratic polynomial over the entire n-dimensional hypercube. Thus polynomial optimization even over this simple-seeming 0/1-polytope is already very interesting and complicated. In max cut the polynomial to be optimized does not only have low degree, but is also symmetric under permutation of variables. One would expect that symmetric polynomial optimization has nice properties since one can bring to bear the power of the group action on the problem.

Greg's talk addressed precisely this situation. Given a symmetric polynomial $p(x)$ of low degree, he discussed the problem of writing it as a sum of squares polynomial modulo the equations of $\{0,1\}^n$. He explained a remarkably simple method to reduce this question to a univariate sum of squares problem on the "levels" of the hypercube which are the numbers $0,\ldots,n$. This is based on the observation that if we define $t := x_1 + \ldots + x_n$, then a symmetric polynomial p in the variables $x_1,\ldots,x_n$ is a polynomial in t.

Using this reduction he was able to derive simple proofs of previously known results such as certain quadratic polynomials can only be equal to sum of squares polynomials of high degree. The method promises to have more applications and Greg gave an excellent exposition of his methods on the black board at the meeting.

4.3 Markus Schweighofer

Markus Schweighofer was the main representative at this seminar from the field of *real algebraic geometry* which is the home of the theory of sums of squares polynomials. The organizers had requested from him a survey talk that would explain the real algebraic methods (via semidefinite programming) for solving polynomial optimization problems. Markus did precisely that. He developed the popular Lasserre hierarchy for solving polynomial optimization problems in an elementary way starting with systems of polynomial inequalities. He showed how the method is a natural consequence of introducing the right kind of combinations of existing inequalities followed by linearizations. This was instructive even to the experts. In the process, Markus introduced the important Positivestellensatz which underlies real algebraic geometry. This is a certificate for the infeasibility of a system of polynomial inequalities analogous to Hilbert's Nullstellensatz that provides a certificate for the infeasibility of systems of polynomial equations. Most of the audience was far from real algebraic geometry and this talk was both a friendly and natural introduction to the theorems in real algebraic geometry that contribute to polynomial optimization and the underlying convex sets. Markus went on to describe recent results which were more advanced, but the talk served well the purpose of providing the audience a taste of real algebraic geometry and its connections to psd rank.

4.4 Alexander Barvinok

In contrast to the setting of extended formulations, Alexander Barvinok talked about approximating a convex body $B \subset \mathbb{R}^d$ by a polytope $P \subset \mathbb{R}^d$ in the same dimensional space. In the work he discussed, based on the paper "Thrifty approximations of convex bodies by polytopes", the goal is to find a polytope P satisfying $P \subset B \subset \tau P$ with as few vertices as possible. He showed that if B is centrally symmetric, then for $\tau = 1 + \epsilon$ for a small constant ϵ one can find a $1 + \epsilon$ approximating polytope with $(\frac{1}{2\sqrt{\epsilon}} \ln(\frac{1}{\epsilon}))^d$ many vertices. This improves by about a square-root factor of ϵ over the standard volumetric argument for constructing an ϵ-net which gives $(1 + 2/\epsilon)^d$ many points. The proof has essentially two steps. The first step relies on John's theorem which says that for a centrally symmetric convex body B, the maximal volume ellipsoid E contained in B satisfies $E \subseteq B \subseteq \sqrt{d}E$. This is used to find a polytope with few vertices (only $O(d)$) that gives a weak $O(\sqrt{d})$ approximation to B. The second step amplifies the quality of the approximation, at the expense of adding more vertices to the approximating polytope. This is done by mapping B to a convex body $\widehat{B}$ in a higher dimensional space where it is argued using Chebyshev polynomials that a weak approximation to $\widehat{B}$ implies a very good approximation of B. Then the previous argument using John's theorem can be applied to $\widehat{B}$.

4.5 James Lee and David Steurer

One of the major motivating questions in the organization of this Dagstuhl seminar was to show superpolynomial lower bounds on the positive semidefinite extension complexity of an explicit polytope. This was achieved by James Lee, David Steurer, and Prasad Raghavendra in November 2014. James and David were in attendance at the seminar and treated us to 3 hours of lectures going through the proof in detail. At a high level, the proof "lifts" lower bounds against the Lasserre/sum of squares hierarchy to lower bounds for positive semidefinite extension complexity. In more detail, for a function $f : \{0,1\}^m \to \mathbb{R}_+$ they consider a matrix with rows indexed by subsets $S \subset [n], |S| = m$ and columns indexed by $x \in \{0,1\}^n$, where n is polynomial in m. The matrix $M_f(S, x) = f(x|_S)$ is shown to have psd rank that is exponential in the sum of squares degree of f.

The basic idea of the proof is the following. If f has sum of squares degree d, then there is a certificate of this fact known as a pseudoexpectation. This is a function that has negative correlation with f, but nonnegative correlation with any sum of squares polynomial of degree $< d$. The proof leverages this pseudodistribution into a functional that separates M_f from the set $\mathcal{S}$ of all matrices with small psd rank and that have the same ℓ_1 and ℓ_∞ norms as M_f. On M_f this functional is negative ($\leq -\epsilon$). However, it is shown that the functional is at least $-\epsilon/2$ on any matrix in $\mathcal{S}$. This is done in two steps. The first step, quantum learning, shows that if N has a small size factorization then the value of the functional on N is close to that of a matrix $\tilde{N}$ that has a factorization $N(S, x) = \mathrm{Tr}(A_S B_x)$ where B_x has bounded sums of squares degree as a matrix polynomial in x. The second step, based on random restrictions, shows that the value of the functional on a matrix with a bounded degree factorization can only be slightly negative $> -\epsilon/2$. Taken together this gives the separating functional as desired.

4.6 Samuel Fiorini

Samuel Fiorini talked about 2-level polytopes. A 2-level polytope is a polytope P where for each facet F there is a parallel facet F' such that $F \cup F'$ contains all vertices of P. Thus the slack matrix of P, after scaling the columns, is a matrix with all entries 0 or 1. As such 2-level polytopes give nice examples of nontrivial boolean matrices of low rank. Fiorini raised the very interesting question of the "log rank conjecture" for 2-level polytopes. This is the question: for the slack matrix of a 2-level polytope is the deterministic communication complexity at most a polynomial in the logarithm of the rank of the matrix? A relaxation of this question mentioned by Fiorini is to show that the extension complexity of any d-dimensional 2-level polytope is at most $2^{\mathrm{polylog(d)}}$. So far the best result known follows from Lovett's results on the log rank conjecture that shows a bound of $2^{\sqrt{d}\log(d)}$.

5 Overview of Talks

5.1 Thrifty approximations of convex bodies by polytopes

Alexander Barvinok (University of Michigan – Ann Arbor, US)

Main reference A. Barvinok, "Thrifty Approximations of Convex Bodies by Polytopes," International Mathematical Research Notices, Vol. 2014(16):4341–4356, 2013; pre-print available as arXiv:1206.3993v2 [math.MG].
URL http://dx.doi.org/10.1093/imrn/rnt078
URL http://arxiv.org/abs/1206.3993v2

This is a survey talk on how well a convex body can be approximated by a polytope with a given number of vertices. We measure the quality of approximation with respect to the Banach-Mazur distance and its versions (that is, given a convex body and an inscribed polytope, by what factor the polytope should be dilated to contain the body) and consider both fine (factors close to 1) and coarse (large factors) approximations.

5.2 Preservers of Completely Positive Matrix Rank

LeRoy B. Beasley (Utah State University, US)

Let $M_{m\times n}(R)$ denote the set of m-by-n matrices with entries in R. We write $M_{m\times n}(R_+)$ to denote the subsets of matrices, all of whose entries are nonnegative. Let $S_n(R)$ denote the set of all n-by-n real symmetric matrices. A matrix $A \in S_n(R)$ is said to be completely positive if there is some matrix $B \in M_{n\times k}(R+)$ such that $A = BB^t$. The CP-rank of the matrix A is the smallest k such that $A = BB^t$ for some $B \in M_{n\times k}(R+)$. In this article we shall investigate the linear operatros on $S_n(R)$ that preserve sets of matrices defined by the CP- rank. We classify those that preserve the CP-rank function, those that preserve the set of CP-rank-1 matrices , those that preserve the sets of CP-rank-1 matrices and the set of CP-rank-2 matrices, and those that strongly preserve the set of CP-rank-1 matrices.

5.3 Symmetry reduction for sums of squares polynomials on the hypercube

G.eg Blekherman (Georgia Institute of Technology – Atlanta, US)

Main reference G. Blekherman, J. Gouveia, J. Pfeiffer, "Sums of Squares on the Hypercube," arXiv:1402.4199v1 [math.AG], 2014.
URL http://arxiv.org/abs/1402.4199v1

Let p be a symmetric polynomial, i.e. a polynomial fixed under permutations of variables, and let H be the discrete hypercube $\{0,1\}^n$. The question of whether p is a sum of squares of polynomials of low degree on H can be reduced to a univariate sum of squares problem. I will present this reduction and explain how some known results on the Lasserre sum of squares hierarchy on H easily follow from it.

5.4 Rescaling PSD factorizations

Daniel Dadush (CWI – Amsterdam, NL)

I will present a short proof that any PSD factorization of a matrix M can be rescaled so that the operator norm of each matrix in the factorization can be bounded by a function of the maximum entry size and rank of the factorization.

5.5 2-level polytopes

Samuel Fiorini (University of Brussels, BE)

A polytope is called 2-level if for each facet-defining hyperplane H, the vertices that are not on H all lie in a hyperplane that is parallel to H. A polytope is 2-level iff it has a binary slack matrix. These polytopes appeared e.g. as a solution of a problem of Lovász in Gouveia, Parrilo and Thomas (2010), and earlier under the name of compressed polytopes in works of Stanley (1980) and Sullivant (2006). In this talk I will give motivations to study 2-level polytopes, survey some results about their structure and state many open problems about them.

5.6 Numerical Computation of Nonnegative and PSD Factorizations

Nicolas Gillis (University of Mons, BE)

Main reference A. Vandaele, N. Gillis, F. Glineur, D. Tuyttens,"Heuristics for Exact Nonnegative Matrix Factorization," arXiv:1411.7245v1 [math.OC], 2014.
URL http://arxiv.org/abs/1411.7245v1

Nonnegative (resp. positive semidefinite–PSD) factorizations allow to compute linear (resp. semidefinite) extended formulations of polytopes. In this talk, we present several numerical algorithms to compute such factorizations using standard low-rank matrix approximation formulations. These algorithms (sometimes) allow to provide explicit extended formulations and hence upper bounds for the nonnegative and PSD ranks (that is, the sizes of the smallest extended formulations). We illustrate this on regular n-gons, and show how our algorithms can give insight on their smallest extended formulations.

Matlab code to compute nonnegative factorizations can be downloaded from https://sites.google.com/site/exactnmf/

5.7 Extension complexity bounds for polygons – Numerical factorizations and conjectures

Francois Glineur (University of Louvain, BE)

Joint work of Glineur, Francois; Nicolas Gillis; Arnaud Vandaele and Julien Dewez

In this talk we describe an explicit nonnegative factorization of the slack matrix of the regular n-gon. We conjecture its rank to be optimal based on extensive numerical computations (cf. Gillis' talk on heuristic exact matrix factorization). In particular we describe the behaviour of the nonnegative tank for values of n lying between powers of two. We also compare this nonnegative rank to available computable lower bounds. Finally we describe our attempts at obtaining a (smaller) explicit positive semidefinite factorization for this slack matrix.

5.8 Positive Semidefinite Rank

João Gouveia (University of Coimbra, PT)

Main reference H. Fawzi, J. Gouveia, P. A. Parrilo, R. Z. Robinson, R. R. Thomas, "Positive semidefinite rank," arXiv:1407.4095v1 [math.OC], 2014.
URL http://arxiv.org/abs/1407.4095v1

In this talk we will cover basic properties of the positive semidefinite rank. A special emphasis will be put on its geometry and complexity, its relation to the square root rank and its space of factorizations. Based on joint work with Hamza Fawzi, Pablo Parrilo, Richard Robinson and Rekha Thomas.

5.9 Linear algebra invariants over semirings

Alexander Guterman (Moscow State University, RU)

Semiring matrix invariants are useful in different problems of linear algebra and its applications. In the talk we will discuss the properties of different matrix invariants which replace rank, determinant, etc., over semirings, in particular, over tropical semirings. The preferences will be given to the interrelations between different invariants and their applications. The talk is based on several works joint with Marianne Akian, LeRoy Beasley, Stephane Gaubert, and Yaroslav Shitov.

5.10 Extended Formulations: Constructions and Obstructions

Volker Kaibel (Universität Magdeburg, DE)

Joint work of Kaibel, Volker; Lee, Jon; Walter, Matthias; Weltge, Stefan

We are going to demonstrate both the power of extended formulations as well as the limitations of the concept. In the first part (based on joint work with Jon Lee, Matthias Walter and Stefan Weltge), we present polynomial size extended formulations for the independence polytopes of regular matroids. On the way, we review beautiful extended formulations for the spanning tree polytopes of planar graphs due to Williams (2001) and encounter some very simple, but stunning occurences of representations in higher dimensional spaces. In the second part (based on joint work with Stefan Weltge), we give an elementary combinatorial proof of an exponential lower bound on the rectangle covering numbers (and hence on the nonnegative ranks) of unique disjointness matrices, and we review the implications for the minimal sizes of extended formulations of the correlation polytopes.

5.11 Semialgebraic geometry of nonnegative and psd ranks

Kaie Kubjas (Aalto Science Institute, FI)

Cohen and Rothblum asked in 1993 whether every rational matrix of nonnegative rank r has a size r rational nonnegative factorization. In joint work with Elina Robeva and Bernd Sturmfels, we answer this question positively for matrices of nonnegative rank 3. I will explain how looking for a semialgebraic description of the set of matrices of nonnegative rank at most 3 helped us to derive this result, and talk about ongoing research with Elina Robeva and Richard Z. Robinson on semialgebraic geometry of the set of matrices of psd rank at most k.

5.12 Lower bounds on semidefinite programming relaxations: Part I

James R. Lee (University of Washington – Seattle, US)

Main reference James R. Lee, Prasad Raghavendra, David Steurer, "Lower bounds on the size of semidefinite programming relaxations," arXiv:1411.6317v1 [cs.CC], 2014.
URL http://arxiv.org/abs/1411.6317v1

We introduce a method for proving lower bounds on the size of SDP relaxations for combinatorial problems. In particular, we show that the cut, TSP, and stable set polytopes on n-vertex graphs are not the affine image of the feasible region of any spectrahedron of dimension less than 2^{n^c} for some constant $c > 0$. A spectrahedron is the feasible region of an SDP: The intersection of the positive semidefinite cone and an affine subspace. This yields the first super-polynomial lower bound on the semidefinite extension complexity of any explicit family of polytopes.

Our results follow from a general technique for proving lower bounds on the positive semidefinite rank of a nonnegative matrix. To this end, we establish a close connection between arbitrary SDPs and those arising from the sum-of-squares hierarchy. For approximating maximum constraint satisfaction problems, we prove that SDPs of polynomial-size are equivalent in power to those arising from degree-O(1) sum-of-squares relaxations. This result implies, for instance, that no family of polynomial-size SDP relaxations can achieve better than a 7/8-approximation for max 3-sat. Part I: After a brief review of the SDP extended formulation model, I will recast the SDP model as a proof system for certifying the nonnegativity of a family of statements. The size of the SDP corresponds to the number of axioms in the system. Then we will attempt to prove that for the family of nonnegative quadratic functions on the discrete cube, the smallest family of axioms is the subspace of low-degree multilinear polynomials. This will be reduced the tak of approximating an arbitrary low-rank PSD factorization by a "low-degree" PSD factorization. In Part II, David will discuss and prove the existence of an approximate low-degree factorization.

5.13 Extension complexity of polytopes with few vertices (or facets)

Arnau Padrol Sureda (FU Berlin, DE)

The number of combinatorial types of d-polytopes with up to d+4 vertices grows superexponentially with d. However, only quadratically many can have realizations with extension complexity smaller than d+4. These are easy to classify into finitely many families and the exact extension complexity of each realization is easy to decide.

5.14 Polytopes of minimum positive semidefinite rank in dimension four

Kanstantsin Pashkovich (University of Waterloo, CA)

For a given polytope the smallest size of a semidefinite extended formulation can be bounded from below by the dimension of the polytope plus one. This talk is about polytopes for which this bound is tight, i.e. polytopes with positive semidefinite (psd) rank equal to their dimension plus one. I will speak about a classification of psd minimum polytopes in dimension four.

Joint work with Gouveia, Robinson and Thomas.

5.15 Completely positive semidefinite cone

Teresa Piovesan (CWI – Amsterdam, NL)

Main reference S. Burgdorf, M. Laurent, T. Piovesan, "On the closure of the completely positive semidefinite cone and linear approximations to quantum colorings," arXiv:1502.02842v1 [math.OC], 2015.
URL http://arxiv.org/abs/1502.02842v1

We investigate structural properties of the completely positive semidefinite cone, consisting of all the n-by-n symmetric matrices that admit a Gram representation by positive semidefinite matrices of any size. This cone has been introduced to model quantum graph parameters as conic optimization problems. Recently it has also been used to characterize the set Q of bipartite quantum correlations, as projection of an affine section of it. We have two main results concerning the structure of the completely positive semidefinite cone, namely about its interior and about its closure. On the one hand we construct a hierarchy of polyhedral cones which covers the interior of the completely positive semidefinite cone, which we use for computing some variants of the quantum chromatic number by way of a linear program. On the other hand we give an explicit description of the closure of the completely positive semidefinite cone, by showing that it consists of all matrices admitting a Gram representation in the tracial ultraproduct of matrix algebras.

Join work with Sabine Burgdorf and Monique Laurent

5.16 Small Linear Programs Cannot Approximate Vertex Cover Within a Factor of $2 - \epsilon$;

Sebastian Pokutta (Georgia Institute of Technology, US)

We prove that every linear programming (LP) relaxation that approximates vertex cover within a factor of $2 - \epsilon$ has super-polynomially many inequalities. As a direct consequence of our methods, we also establish that LP relaxations that approximate independent set within any constant factor have super-polynomially many inequalities.

5.17 Geometric properties, matroids, and forbidden minors

Raman Sanyal (FU Berlin, DE)

Nonegative rank, PSD-rank, and Theta rank are just some geometric properties that are interesting but difficult to compute. In this talk I will discuss such geometric properties for matroid base polytopes. In this setup, the class of matroids with bounded rank is closed under taking minors. I will give some excluded-minor descriptions and connections to classical results in matroid theory. This is joint work with Francesco Grande.

5.18 On the exactness of moment relaxations

Markus Schweighofer (Universität Konstanz, DE)

Semidefinite programming is optimizing a linear function over a spectrahedron, i.e., the solution set of a linear matrix inequality. It becomes increasingly important in discrete and continuous optimization since it serves to solve the hierarchy of moment relaxations for polynomial optimization problems introduced by Lasserre around the turn of the millennium. In our days, algorithms like the Goemans-Williamson relaxation for the maximum cut problem can be viewed as the first level of this hierarchy. Recent work of Lee, Raghavendra and Steurer seems to suggest that moment relaxations might be universal semidefinite programming relaxations. The semidefinite programs dual to moment relaxations are sum-of-squares programs. In this talk, we will try to analyze the moment relaxations for writing a convex semi-algebraic set as a projected spectrahedron. We have a new positive result building upon the work of Helton and Nie.

5.19 Sublinear extensions of polygons

Yaroslav Shitov (NRU Higher School of Economics – Moscow, RU)

Main reference Y. Shitov, "Sublinear extensions of polygons," arXiv:1412.0728v1 [math.CO], 2014.
URL http://arxiv.org/abs/1412.0728v1

One of the central questions of extended formulations theory is the existence of strong lower bounds for polytopes naturally arising in optimization. The study of lower bounds for linear formulations dates back to the 1991 paper by Yannakakis, and, since then, many important polytopes have been proven to have exponential extension complexity.

This talk will be devoted to the case of generic polytopes, that is, polytopes whose vertices chosen randomly in space. How strong are linear formulations in this case? Beasley and Laffey conjectured that, even for a convex n-gon on the plane, the extension complexity can be worst possible (equal to n). However, it turns out that linear formulations are not that weak: we show that the extension complexity of any convex n-gon is o-small of n. Moreover, we can provide examples of generic n-gons whose complexities do not exceed $c \cdot \sqrt{n}$; this upper bound coincides with the known lower bound up to a constant factor. Our results allow us to make a number of conjectures concerning higher-dimensional generic polytopes.

5.20 Lower bounds for semidefinite programming relaxations: part II

David Steurer (Cornell University – Ithaca, US)

Main reference J. R. Lee, P. Raghavendra, D. Steurer, "Lower bounds on the size of semidefinite programming relaxations," arXiv:1411.6317v1 [cs.CC], 2014.
URL http://arxiv.org/abs/1411.6317v1

We introduce a method for proving lower bounds on the size of SDP relaxations for combinatorial problems. In particular, we show that the cut, TSP, and stable set polytopes on n-vertex graphs are not the affine image of the feasible region of any spectrahedron of dimension less than 2^{n^c} for some constant $c > 0$. A spectrahedron is the feasible region of an SDP: The intersection of the positive semidefinite cone and an affine subspace. This yields the first super-polynomial lower bound on the semidefinite extension complexity of any explicit family of polytopes.

Our results follow from a general technique for proving lower bounds on the positive semidefinite rank of a nonnegative matrix. To this end, we establish a close connection between arbitrary SDPs and those arising from the sum-of-squares hierarchy. For approximating maximum constraint satisfaction problems, we prove that SDPs of polynomial-size are equivalent in power to those arising from degree-O(1) sum-of-squares relaxations. This result implies, for instance, that no family of polynomial-size SDP relaxations can achieve better than a 7/8-approximation for max 3-sat.

Part II: We will show that for certain families of matrices, every low-rank PSD factorization can be approximated by a low-degree PSD factorization. Following Part I, this will reduce lower bounds for PSD rank to lower bounds for sum-of-squares degree.

5.21 Subgraph Polytopes and Independence Polytopes of Count Matroids

Stefan Weltge (Universität Magdeburg, DE)

Given an undirected graph, the non-empty subgraph polytope is the convex hull of the characteristic vectors of pairs (F,S) where S is a non-empty subset of nodes and F is a subset of the edges with both endnodes in S. We obtain a strong relationship between the non-empty subgraph polytope and the spanning forest polytope. We further show that these polytopes provide polynomial size extended formulations for independence polytopes of count matroids, which generalizes recent results obtained by Iwata et al. referring to sparsity matroids. As a byproduct, we obtain new lower bounds on the extension complexity of the spanning forest polytope in terms of extension complexities of independence polytopes of these matroids.

5.22 Quantum Communication Complexity as a Tool to Analyze PSD Rank

Ronald de Wolf (CWI – Amsterdam, NL)

Main reference Jedrzej Kaniewski, Troy Lee, Ronald de Wolf, "Query complexity in expectation," arXiv:1411.7280v1 [quant-ph], 2014.
URL http://arxiv.org/abs/1411.7280v1

We start with a brief introduction to the general area of quantum communication complexity, and then describe the connection with the positive-semidefinite rank of matrices (due to Fiorini et al. '12): the logarithm of the psd rank of a matrix M equals the minimal quantum communication needed by protocols that compute M in expectation. This means that results about quantum communication complexity can be used to obtain results about psd rank, both lower and upper bounds. As an example of the latter we present an efficient quantum communication protocol (due to Kaniewski, Lee, and de Wolf '14) that induces an exponentially-close approximation for the slack matrix for the perfect matching polytope, of psd rank only roughly $2^{\sqrt{n}}$. In contrast, Braun and Pokutta'14 showed that such approximating matrices need nonnegative rank $2^{\Omega(n)}$.

6 Problems resolved from Dagstuhl Seminar 13082

Progress has been made on a number of open problems discussed at the preceding seminar in February 2013. We briefly highlight some of this work here.

1. Is there a nonnegative n-by-n matrix of rank 3 and nonnegative rank n (for $n \geq 7$)? Shortly after the last seminar Yaroslav Shitov showed that this is not the case. He has since improved these results to show that every nonnegative matrix of rank 3 has nonnegative rank $o(n)$ in this paper http://arxiv.org/abs/1412.0728.
2. Does $\text{rank}_+(A \otimes B) = \text{rank}_+(A) \cdot \text{rank}_+(B)$? This was disproven by Hamza Fawzi using the software to compute nonnegative rank of Nicolas Gillis. The analogous statement for positive semidefinite rank is also false as shown by Lee, Wei, and de Wolf http://arxiv.org/abs/1407.4308.
3. Does positive semidefinite rank depend on the underlying field? Say that a matrix M has rational entries, and we require the matrices A_i, B_j in the factorization $M_{ij} = \text{Tr}(A_i B_j)$ to have rational entries as well. Can this require larger matrices than if we allowed A_i, B_j to have real entries. Gouveia, Fawzi, and Robinson indeed give an example of a rational matrix M where the rational psd rank is larger than the real psd rank http://arxiv.org/abs/1404.4864. Lee, Wei, and de Wolf give an example of a family of matrices where the real psd rank is asymptotically larger than the complex psd rank (where the factors are allowed to be complex Hermitian psd matrices) by a factor of $\sqrt{2}$ http://arxiv.org/abs/1407.4308.
4. What is the largest possible gap between the approximate rank and the approximate nonnegative rank for boolean matrices? Let $\text{rank}_{1/10}(M)$ denote the minimum rank of a matrix S such that $\|S - M\|_\infty \leq 1/10$. Define $\text{rank}^+_{1/10}(M)$ in the same way for the nonnegative rank. Kol, Moran, Shpilka, and Yehudayoff (Approximate Nonnegative Rank is Equivalent to the Smooth Rectangle Bound, ICALP 2014) made progress on

this question by showing that the approximate nonnegative rank of the set intersection matrix of subsets of $[n]$ is $2^{\Omega(n)}$ while it is known that its approximate rank is $2^{O(\sqrt{n})}$.

5. For a n-by-n boolean matrix A satisfying $A \circ A^T = I_n$ how small can the rank be? Here $\circ$ denotes the entrywise product of matrices A^T is the transpose of A. This question is motivated by the fooling set method in communication complexity. It is known that the rank must be at least $\sqrt{n}$. Friesen, Hamed, Lee, and Theis http://arxiv.org/abs/1208.2920 constructed a growing family of matrices with rank $\sqrt{(n)} + O(1)$ over a finite field, and another construction of matrices with integer entries with rank $O(\sqrt{n})$ over the reals. Shigeta and Amano (http://arxiv.org/abs/1311.6192) have now essentially answered the original question, constructing a growing family of *boolean* matrices with rank $n^{1/2+o(1)}$.

7 Open Problems

Many open problems were posed during the talks at the seminar. There was also a session on Wednesday evening for presenting open problems at the board and discussing them with the participants. Jon Swenson from the University of Washington kindly recorded several of these problems on behalf of the organizers. The main questions that came up are listed below. Some of these problems are classical—the name indicates the promotor of the problem, not necessarily the originator.

1. (R. de Wolf) For a regular n-gon it is known that the nonnegative rank of the slack matrix is $\Theta(\log(n)$. For the psd rank, the best lower bound is $\Omega(\sqrt{\frac{\log n}{\log \log n}})$ and the best upper bound is still $O(\log n)$. What is the true psd rank of the slack matrix of the n-gon?
2. (T. Lee) Let the *rational* psd rank of a matrix $A \in \mathbb{R}_+^{m \times n}$ be the minimum r such that $A(i,j) = \frac{B(i,j)}{C(i,j)}$ for all $i \in [m], j \in [n]$ where B, C are nonnegative matrices with $r = \text{rank}_{psd}(B) + \text{rank}_{psd}(C)$. Does the slack matrix of the correlation polytope have exponential rational psd rank? Note for contrast that the slack matrix of the perfect matching polytope has rational psd rank $O(n^4)$.
3. (J. Gouveia) Let $\mathcal{F}_k(M) = \mathcal{SF}(M)/\text{GL}(k)$ be the space of factorizations of a matrix M with psd rank k. Under what conditions is $\mathcal{F}_k(M)$ connected? Perhaps one can come up with a condition involving psd rank, ordinary rank, or something else. (For notation, see http://www.mat.uc.pt/~jgouveia/dagstuhl.pdf.)
4. (K. Pashkovich) Give a geometric characterization of the polytopes in $\mathbb{R}^n$ for which the psd rank of the slack matrix of P is equal to $\dim(P) + 1$. This is the smallest that the psd rank of a polytope can be. Such a characterization exists now for $n \leq 4$.
5. (S. Fiorini) Is there a constant k such that for every d and 2-level polytope P of dimension d the linear extension complexity is $2^{O(\log(d)^k)}$?
6. (V. Kaibel) From results of Rothvoß we know that there is a matroid polytope with exponential linear extension complexity. Give an explicit example of an explicit matroid polytope with exponential linear extension complexity.
7. (J. Lee) What is the positive semidefinite rank of the matching polytope?
8. (J. Lee) We currently know that the slack matrix of the correlation polytope $\text{CORR}_n(x,y) =$ convex hull($x^T y$ for $x, y \in \{0,1\}^n$) satisfies $N^{2/13} \leq \log_2(\text{rank}_{\text{psd}}(\text{CORR}_n)) \leq N$ where $N = 2^n$. What is the right answer? This is the same as the rank of M(f,x)=f(x) where f ranges over all quadratic nonlinear functions that are nonnegative on the hypercube.

9. (J. Lee) Consider the matrix $M(G, x) = 0.99 - val_G(x)$ where G ranges over n-vertex graphs, x ranges over cuts, and $val_G(x)$ is the fraction of edges in G that cross the cut x. What is the non-negative rank of M? The best lower bound we have is $n^{\Omega(logn)}$, and there is a barrier of sorts because the n^{-10}-approximate nonnegative rank (additive error) is $n^{O(\log n)}$.
10. (J. Lee) Can one exhibit an explicit function $f : \{0,1\}^n \to \mathbb{R}_+$ so that f cannot be written as a sum of squares of s-sparse polynomials with $s = n^{O(1)}$? Here s-sparse means that the polynomial contains at most s monomials.

8 Problems resolved from Dagstuhl Seminar 15082

Cohen and Rothblum raised the question in 1993 if the nonnegative rank of a rational matrix over the reals is the same as its nonnnegative rank over the rationals. More generally they asked: if A has entries in a sub-semiring S of the semiring R of the nonnegative reals, how large can the gap be between $\text{rank}_S(A)$ and $\text{rank}_R(A)$? For the latter question, Yaroslav Shitov (http://arxiv.org/abs/1505.01893) has recently provided the first example where $\text{rank}_S(A) \neq \text{rank}_R(A)$.

Participants

- Alexander Barvinok
University of Michigan – Ann Arbor, US
- LeRoy B. Beasley
Utah State University, US
- Greg Blekherman
Georgia Inst. of Technology, US
- Daniel Dadush
CWI – Amsterdam, NL
- Ronald de Wolf
CWI – Amsterdam, NL
- Samuel Fiorini
University of Brussels, BE
- Nicolas Gillis
University of Mons, BE
- Francois Glineur
University of Louvain, BE
- João Gouveia
University of Coimbra, PT
- Sebastian Gruler
Universität Konstanz, DE
- Alexander Guterman
Moscow State University, RU
- Volker Kaibel
Universität Magdeburg, DE
- Hartmut Klauck
Nanyang TU – Singapore, SG
- Kaie Kubjas
Aalto Science Institute, FI
- James R. Lee
University of Washington – Seattle, US
- Troy Lee
National Univ. of Singapore, SG
- Arnau Padrol Sureda
FU Berlin, DE
- Kanstantsin Pashkovich
University of Waterloo, CA
- Teresa Piovesan
CWI – Amsterdam, NL
- Sebastian Pokutta
Georgia Inst. of Technology, US
- Raman Sanyal
FU Berlin, DE
- Markus Schweighofer
Universität Konstanz, DE
- Yaroslav Shitov
NRU Higher School of Economics – Moscow, RU
- David Steurer
Cornell University – Ithaca, US
- Jonathan Swenson
University of Washington – Seattle, US
- Dirk Oliver Theis
University of Tartu, EE
- Rekha R. Thomas
University of Washington – Seattle, US
- Hans Raj Tiwary
Charles University – Prague, CZ
- Antonios Varvitsiotis
National Univ. of Singapore, SG
- Stefan Weltge
Universität Magdeburg, DE