

Exponential Lower Bounds for Polytopes in Combinatorial Optimization

Samuel Fiorini^{*1}, Serge Massar^{†3}, Sebastian Pokutta², Hans Raj Tiwary^{‡1}, and Ronald de Wolf^{§4}

¹Department of Mathematics, Université libre de Bruxelles CP 216, Boulevard du Triomphe, 1050 Brussels, Belgium. *Email:* {sfiorini, htiwary}@ulb.ac.be

²Industrial and Systems Engineering (ISyE) and Algorithms and Randomness Center (ARC), Georgia Institute of Technology, Groseclose 0205, Atlanta, GA 30332, USA. *Email:* sebastian.pokutta@isye.gatech.edu

³Laboratoire d'Information Quantique, Université libre de Bruxelles CP 225, Boulevard du Triomphe, 1050 Brussels, Belgium. *Email:* smassar@ulb.ac.be

⁴CWI and University of Amsterdam, Amsterdam, The Netherlands. *Email:* rdewolf@cwi.nl

December 24, 2013

Abstract

We solve a 20-year old problem posed by Yannakakis and prove that there exists no polynomial-size linear program (LP) whose associated polytope projects to the traveling salesman polytope, even if the LP is not required to be symmetric. Moreover, we prove that this holds also for the cut polytope and the stable set polytope. These results were discovered through a new connection that we make between one-way quantum communication protocols and semidefinite programming reformulations of LPs.

1 Introduction

Since the advent of the simplex method [20], linear programming has become a prominent tool for solving optimization problems in practice. On the theoretical side, LPs can be solved in polynomial time via either the ellipsoid method [44] or interior point methods [42].

In 1986–1987 there were attempts [68] to prove $P = NP$ by giving a polynomial-size LP that would solve the traveling salesman problem (TSP). Due to the large size and complicated structure of the proposed LP for the TSP, it was difficult to show directly that the LP was erroneous. In a groundbreaking effort to refute all such attempts, Yannakakis [71] proved that every symmetric LP for the TSP has exponential size (see [72] for the journal version). Here, an LP is called *symmetric* if every permutation of the cities can be extended to a permutation of all the variables of the LP that preserves the constraints of the LP. Because the proposed LP for the TSP was symmetric, it could not possibly be correct.

In his paper, Yannakakis left as a main open problem the question of proving that the TSP admits no polynomial-size LP, *symmetric or not*. We solve this question by proving a super-polynomial lower bound on the number of inequalities in *every* LP for the TSP. We also prove such unconditional super-polynomial lower bounds for the maximum cut and maximum stable set problems. Therefore, it is impossible to prove $P = NP$ by means of a polynomial-size LP that expresses any of these problems. Our approach is inspired by a close

^{*}Partially supported by the *Actions de Recherche Concertées* (ARC) fund of the French community of Belgium.

[†]Partially supported by the European Commission under the project QCS (Grant No. 255961).

[‡]Postdoctoral Researcher of the *Fonds National de la Recherche Scientifique* (F.R.S.–FNRS).

[§]Partially supported by a Vidi grant from the Netherlands Organization for Scientific Research (NWO), and by the European Commission under the project QCS (Grant No. 255961).

connection between semidefinite programming reformulations of LPs and one-way quantum communication protocols that we introduce here.

1.1 State of the Art

From Problems to Polytopes For combinatorial optimization problems such as the TSP, the feasible solutions can be encoded as points in a set $X \subseteq \{0, 1\}^d$ in such a way that solving an instance of the problem amounts to optimizing a linear objective function over X , with coefficients given by the instance. By taking the convex hull of X , one obtains a polytope $P := \text{conv } X$ (see Appendix A for background on polytopes). Optimizing any linear function $f(x)$ over X is equivalent to optimizing this function $f(x)$ over $P = \text{conv } X$.

For example, for the TSP we have a set $X \subseteq \{0, 1\}^{\binom{n}{2}}$ of 0/1-points that correspond to a Hamiltonian cycle in the complete n -vertex graph K_n . The convex hull of these points is the *TSP polytope* $\text{TSP}(n) = \text{conv } X$. An instance of the TSP is given by the set of edge-weights w_{ij} . Solving this instance amounts to minimizing $f(x) := \sum_{i < j} w_{ij} x_{ij}$ over all $x \in \text{TSP}(n)$. This minimum is attained at a vertex of the polytope, i.e., at a point $x \in X$.

The idea of representing the set of feasible solutions of a problem by a polytope forms the basis of a standard and powerful methodology in combinatorial optimization, see, e.g., [64].

Extended Formulations and Extensions Resuming the discussion above (and assuming that the problem is a minimization problem), we have $\min\{f(x) \mid x \in X\} = \min\{f(x) \mid x \in P\} = \min\{f(x) \mid Ax \leq b\}$, where $Ax \leq b$ is any linear description of P . This turns any given instance of the combinatorial optimization problem into an LP, however over an implicit system of constraints the LP is potentially large since it has at least one inequality per facet of P . In fact, even for polynomially solvable problems, the associated polytope P may have an exponential number of facets.

By working in an extended space, that is, considering extra variables $y \in \mathbb{R}^r$ besides the original variables $x \in \mathbb{R}^d$, it is often possible to decrease the number of constraints. In some cases, a polynomial increase in dimension can be traded for an exponential decrease in the number of constraints. This is the idea underlying extended formulations.

Formally, an *extended formulation* (EF) of a polytope $P \subseteq \mathbb{R}^d$ is a linear system

$$E^=x + F^=y = g^=, \quad E^{\leq}x + F^{\leq}y \leq g^{\leq} \tag{1}$$

in variables $(x, y) \in \mathbb{R}^{d+r}$, where $E^=, F^=, E^{\leq}, F^{\leq}$ are real matrices with d, r columns respectively, and $g^=, g^{\leq}$ are column vectors, such that $x \in P$ if and only if there exists y such that (1) holds.

The *size* of an EF is defined as the number of *inequalities* in the system. Another possible definition of size would be the sum of the number of variables and total number of constraints (equalities plus inequalities) defining the EF. This would make little difference because if a polytope $P \subseteq \mathbb{R}^d$ has an EF with r inequalities, then it has an EF with $d+r$ variables, r inequalities and at most $d+r$ equalities (see Remark 1 below for a proof). If we assume that P is full-dimensional (otherwise one may cheat and make d artificially high) then $d \leq r$ and thus the two measures of size are within a constant of each other.

Notice that optimizing any (not necessarily linear) objective function $f(x)$ over all $x \in P$ amounts to optimizing $f(x)$ over all $(x, y) \in \mathbb{R}^{d+r}$ satisfying (1), provided (1) defines an EF of P .

Here, we often restrict to EFs in slack form, that is, containing only equalities and one nonnegativity inequality per additional variable:

$$Ex + Fy = g, \quad y \geq \mathbf{0}. \tag{2}$$

The proof of the factorization theorem (Theorem 3) shows that this can be done without loss of the generality, see Remark 1. In the following we put EFs in slack form to ease the generalization to arbitrary cones. Notice that the size of an EF in slack form can equivalently be defined as the number of additional variables since the only inequalities are from $y \geq \mathbf{0}$.

An *extension* of the polytope P is another polytope $Q \subseteq \mathbb{R}^e$ such that P is the image of Q under a linear map. We define the *size* of an extension Q as the number of facets of Q . If P has an extension of size r ,

then it has an EF of size r . Conversely, it is known that if P has an EF of size r , then it has an extension of size at most r (see Theorem 3). In this sense, the concepts of EF and extension are equivalent.

The Impact of Extended Formulations EFs have pervaded the areas of discrete optimization and approximation algorithms for a long time. For instance, Balas’s disjunctive programming [6], the Sherali-Adams hierarchy [66], the Lovász-Schrijver closures [55], lift-and-project [7], and configuration LPs are all based on the idea of working in an extended space. Recent surveys on EFs in the context of combinatorial optimization and integer programming are [19, 69, 39, 70].

Symmetry Matters Yannakakis [72] proved a $2^{\Omega(n)}$ lower bound on the size of any *symmetric* EF of the TSP polytope $\text{TSP}(n)$ (defined above and in Section 3.4). Although he remarked that he did “not think that asymmetry helps much”, it was recently shown by Kaibel et al. [40] (see also [58]) that symmetry is a restriction in the sense that there exist polytopes that have polynomial-size EFs but no polynomial-size symmetric EF. This revived Yannakakis’s tantalizing question about unconditional lower bounds. That is, bounds which apply to the *extension complexity* of a polytope P , defined as the minimum size of an EF of P (irrespective of any symmetry assumption).

0/1-Polytopes with Large Extension Complexity The strongest unconditional lower bounds so far were obtained by Rothvoß[61]. By an elegant counting argument inspired by Shannon’s theorem [65], it was proved that there exist 0/1-polytopes in \mathbb{R}^d whose extension complexity is at least $2^{d/2-o(d)}$. However, Rothvoß’s counting technique does not provide *explicit* 0/1-polytopes with an exponential extension complexity.

The Factorization Theorem Yannakakis [72] discovered that the extension complexity of a polytope P is determined by certain factorizations of an associated matrix, called the *slack matrix* of P , that records for each pair (F, v) of a facet F and vertex v , the algebraic distance of v to a valid hyperplane supporting F . Defining the *nonnegative rank* of a matrix M as the smallest natural number r such that M can be expressed as $M = TU$ where T and U are nonnegative matrices (i.e., matrices whose elements are all nonnegative) with r columns (in case of T) and r rows (in case of U), respectively, it turns out that the extension complexity of every polytope P is exactly the nonnegative rank of its slack matrix.

We point out that this result generalizes to *any* slack matrix of the polytope, which may contain additional rows corresponding to faces F of P which are not facets and/or additional columns corresponding to points v of P that are not vertices. This fact is used in the proof of our lower bounds on extension complexity, starting with Theorem 7.

This *factorization theorem* led Yannakakis to explore connections between EFs and communication complexity. Let S denote the slack matrix of the polytope P . He proved that: (i) every deterministic communication protocol of complexity k computing S gives rise to an EF of P of size at most 2^k , provided S is a 0/1-matrix; (ii) the nondeterministic communication complexity of the *support matrix* of S (i.e., the binary matrix that has 0-entries exactly where S is 0) yields a lower bound on (the base-2 logarithm¹ of) the extension complexity of P , or more generally, the nondeterministic communication complexity of the support matrix of every nonnegative matrix M yields a lower bound on (the base-2 logarithm of) the nonnegative rank of M .²

Tighter Communication Complexity Connection Faenza et al. [26] proved that the base-2 logarithm of the nonnegative rank of a matrix equals, up to a small additive constant, the minimum complexity of a randomized communication protocol *with nonnegative outputs* that computes the matrix *in expectation*.

¹All logarithms in this paper are in base 2.

²The classical nondeterministic communication complexity of a binary communication matrix is defined as $\lceil \log B \rceil$, where B is the minimum number of monochromatic 1-rectangles that cover the matrix, see [47]. This last quantity is also known as the *rectangle covering bound*. It is easy to see that the rectangle covering bound of the support matrix of any matrix M lower bounds the nonnegative rank of M (see Theorem 4 below).

In particular, every EF of size r can be regarded as such a protocol of complexity $\log r + O(1)$ bits that computes a slack matrix in expectation.

The Clique vs. Stable Set Problem When P is the stable set polytope $\text{STAB}(G)$ of a graph G (see Section 3.3), the slack matrix of P contains an interesting row-induced 0/1-submatrix that is the communication matrix of the *clique vs. stable set problem* (also known as the *clique vs. independent set problem*): its rows correspond to the cliques and its columns to the stable sets (or independent sets) of graph G . The entry for a clique K and stable set S equals $1 - |K \cap S|$. Yannakakis [72] gave an $O(\log^2 n)$ deterministic protocol for the clique vs. stable set problem, where n denotes the number of vertices of G . This gives a $2^{O(\log^2 n)} = n^{O(\log n)}$ size EF for $\text{STAB}(G)$ whenever the whole slack matrix is 0/1, that is, whenever G is a perfect graph.

A notoriously hard open question is to determine the communication complexity (in the deterministic or nondeterministic sense) of the clique vs. stable set problem. (For recent results that explain why this question is hard, see [48, 49].) The best lower bound to this day is due to [37]: they obtained a $\frac{6}{5} \log n - O(1)$ lower bound. Furthermore, they state a graph-theoretical conjecture that, if true, would imply a $\Omega(\log^2 n)$ lower bound, and hence settle the communication complexity of the clique vs. stable set problem. Moreover it would give a worst-case $n^{\Omega(\log n)}$ lower bound on the extension complexity of stable set polytopes. However, a solution to the Huang-Sudakov conjecture seems far away.

1.2 Our Contribution

Our contribution in this paper is three-fold.

- First, inspired by earlier work [23], we define a $2^n \times 2^n$ matrix $M = M(n)$ and show that the nonnegative rank of M is $2^{\Omega(n)}$ because the nondeterministic communication complexity of its support matrix is $\Omega(n)$. The latter was proved by de Wolf [23] using the well-known disjointness lower bound of Razborov [60]. We use the matrix M to prove a $2^{\Omega(n)}$ lower bound on the extension complexity of the cut polytope $\text{CUT}(n)$ (Section 3.2). That is, we prove that *every* EF of the cut polytope has an exponential number of inequalities. Via reductions, we infer from this: (i) an infinite family of graphs G such that the extension complexity of the corresponding stable set polytope $\text{STAB}(G)$ is $2^{\Omega(n^{1/2})}$, where n denotes the number of vertices of G (Section 3.3); (ii) that the extension complexity of the TSP polytope $\text{TSP}(n)$ is $2^{\Omega(n^{1/2})}$ (Section 3.4).

In addition to simultaneously settling the above-mentioned open problems of Yannakakis [72] and Rothvoß[61], our results provide a lower bound on the extension complexity of stable set polytopes that goes much beyond what is implied by the Huang-Sudakov conjecture (thanks to the fact that we consider a different part of the slack matrix). Although our lower bounds are strong, unconditional and apply to explicit polytopes that are well-known in combinatorial optimization, they have very accessible proofs.

- Second, we generalize the factorization theorem to *conic* EFs, that is, reformulations of an LP through a conic program. In particular, this implies a factorization theorem for *semidefinite* EFs: the *semidefinite extension complexity* of a polytope equals the *positive semidefinite rank* (*PSD rank*) of its slack matrix.
- Third, we generalize the tight connection between linear³ EFs and classical communication complexity found by Faenza et al. [26] to a tight connection between semidefinite EFs and quantum communication complexity.⁴ We show that any *rank- r PSD factorization* of a (nonnegative) matrix M gives rise to a one-way quantum protocol computing M in expectation that uses $\log r + O(1)$ qubits and, *vice versa*, that any one-way quantum protocol computing M in expectation that uses q qubits results in a PSD

³In this paragraph, and later in Sections 4 and 5, an EF (in the sense of the previous section) is called a *linear* EF. The use of adjectives such as “linear”, “semidefinite” or “conic” will help distinguishing the different types of EFs.

⁴After a first version of this paper appeared, [38, Theorem 2] have used this notion of PSD rank to characterize the number of qubits of communication between Alice and Bob needed to generate a shared probability distribution.

factorization of M of rank 2^q . Via the semidefinite factorization theorem, this yields a characterization of the semidefinite extension complexity of a polytope in terms of the minimum complexity of (one-way) quantum protocols that compute the corresponding slack matrix in expectation.

Then, we give a complexity $\log r + O(1)$ quantum protocol for computing a nonnegative matrix M in expectation, whenever there exists a rank- r matrix N such that M is the entry-wise square of N . This implies in particular that every d -dimensional polytope with 0/1 slacks has a semidefinite EF of size $O(d)$.

Finally, we obtain an exponential separation between classical and quantum protocols that compute our specific matrix $M = M(n)$ in expectation. On the one hand, our quantum protocol gives a rank- $O(n)$ PSD factorization of M . On the other hand, the nonnegative rank of M is $2^{\Omega(n)}$ because the nondeterministic communication complexity of the support matrix of M is $\Omega(n)$. Thus we also obtain an exponential separation between PSD rank and nonnegative rank.

We would like to point out that the lower bounds on the extension complexity of polytopes established in Section 3 were obtained by first finding an efficient PSD factorization or, equivalently, an efficient one-way quantum communication protocol for the matrix $M = M(n)$. In this sense our classical lower bounds stem from quantum considerations somewhat similar in style to [43, 1, 2]. See [25] for a survey of this line of work.

We would also like to point out that some of our results in Sections 4 and 5 were also obtained by Gouveia, Parrilo and Thomas. This applies to Theorem 14, Corollary 16, Theorem 19 and Corollary 20. We were aware of the fact that they had obtained Theorem 14 and Corollary 16 prior to writing this paper. However, their proofs were not yet publicly available at that time. Theorem 19 and Corollary 20 were obtained independently, and in a different context. All their results are now publicly available, see [35]. Also, after a preprint of our paper had appeared, we learned that Klauck et al. [46] had independently found a matrix (similar but not quite the same as ours) with an exponential separation between PSD rank and nonnegative rank.

1.3 Other Related and Subsequent Work

Yannakakis’s paper has deeply influenced the TCS community. In addition to the works cited above, it has inspired a whole series of papers on the quality of restricted *approximate* EFs, such as those defined by the Sherali-Adams hierarchies and Lovász-Schrijver closures starting with [3] ([4] for the journal version), see, e.g., [16, 63, 28, 18, 32, 31, 8].

After the conference version of our paper appeared, there has been a lot of follow-up work, including on approximations. Braun et al. [10] developed a general framework for studying the power of approximate EFs, independent of specific hierarchies. In particular, via lower bounds on the extension complexity of approximations of the cut polytope, they showed that linear programs for approximating Max-Clique to within a factor $n^{1/2-\epsilon}$ need size at least $2^{\Omega(n^\epsilon)}$. Similarly, they show the existence of a spectrahedron of small size that cannot be approximated by any LP with a polynomial number of inequalities within a factor of $n^{1/2-\epsilon}$. Braverman and Moitra [13] used methods from information complexity to show the same size lower bound even for approximation factor $n^{1-\epsilon}$; Braun and Pokutta [12] subsequently simplified and generalized their result and Braun et al. [11] show that the amortized log nonnegative rank is characterized by information. Such inapproximability results should be contrasted with Håstad’s famous result [36] that it is hard to approximate Max-Clique to within a factor $n^{1-\epsilon}$: Håstad’s result is more general (it is a lower bound for *all* algorithms rather than just linear programs based on extended formulations) but it is conditional on the unproven assumption that $\text{RP} \neq \text{NP}$.

Braun, Fiorini and Pokutta [9] analyze the average-case polyhedral complexity of the maximum stable set problem showing that the extension complexity of the stable set polytope is high for almost all graphs. Pokutta and Van Vyve [59] proved lower bounds on extension complexity for the knapsack problem, and Avis and Tiwary [5] proved lower bounds for the subset-sum and three-dimensional matching problems, as well as others. Kaibel and Weltge [41] gave a more direct proof of the lower bound for the cut polytope, via bounding the size of the largest rectangle in the slack matrix, however, they still use the same set of 2^n valid constraints that we use here (Lemma 6).

Chan et al. [17] prove super-polynomial lower bounds on approximate EFs for MAX CSPs (constraint satisfaction problems). In particular, they prove that every $(2 - \varepsilon)$ -approximate (linear) EF for Max-Cut has $n^{\Omega(\frac{\log n}{\log \log n})}$ size. This is striking because the celebrated approximation algorithm of Goemans and Williamson [33] is based on a $\Theta(n)$ -size semidefinite EF with an approximation factor of at most 1.14. Again, the result of Chan et al. [17] on Max-Cut matches the algorithmic hardness of the problem Khot et al. [45], which assumes the Unique Games Conjecture.

Rothvoß[62] proves that the matching polytope has extension complexity $2^{\Omega(n)}$, solving the second part of the main open problem in [72]. This is the first time such a strong bound is obtained for a polytope over which one can optimize in polynomial time. Rothvoß’s groundbreaking result implies in particular that the extension complexity of the TSP polytope is $2^{\Omega(n)}$, thus going beyond our $2^{\Omega(n^{1/2})}$ lower bound.

Not much is known yet about lower bounds on *semidefinite* EFs. Extending the work of Rothvoß[61], Briët, Dadush and Pokutta [14] show that most 0/1 polytopes (i.e., polytopes that are the convex hull of a random subset of $\{0, 1\}^d$) need exponentially large semidefinite EFs. Fawzi and Parrilo [27] give exponential lower bounds on the size of semidefinite EFs of explicit polytopes in a restricted setting, where the underlying cone is not the full PSD cone but rather a product of fixed-dimensional PSD cones. Lee and Theis [50] obtain polynomial lower bounds based on the support pattern of slack matrices.

Finally, Fiorini et al. [30] use the notion of conic extensions and its relation to communication complexity to study generalized probabilistic theories, which are different from the usual classical or quantum-mechanical theories, and show that all polynomially-definable 0/1-polytopes have small extension complexity with respect to the completely positive cone.

1.4 Organization

The discovery of our lower bounds on extension complexity crucially relied on finding the right matrix M and the right polytope whose slack matrix contains M . In our case, we found these through a connection with quantum communication. However, these quantum aspects are not strictly necessary for the resulting lower bound proof itself. Hence, in order to make the main results more accessible to those without background or interest in quantum computing, we start by giving a purely classical presentation of those lower bounds.

In Section 2 we define our matrix M and lower bound the nondeterministic communication complexity of its support matrix. In Section 3 we embed M in the slack matrix of the cut polytope in order to lower bound its extension complexity; further reductions then give lower bounds on the extension complexities of the stable set, and TSP polytopes. In Section 4 we state and prove the factorization theorem for arbitrary closed convex cones. In Section 5 we establish the equivalence of PSD factorizations of a (nonnegative) matrix M and one-way quantum protocols that compute M in expectation, and give an efficient quantum protocol in the case where some entry-wise square root of M has small rank. This is then used to provide an exponential separation between quantum and classical protocols for computing a matrix in expectation (equivalently, an exponential separation between nonnegative rank and PSD rank). Concluding remarks are given in Section 6.

2 A Simple Matrix with Large Rectangle Covering Bound

In this section we consider the following $2^n \times 2^n$ matrix $M = M(n)$ with rows and columns indexed by n -bit strings a and b , and real nonnegative entries:

$$M_{ab} := (1 - a^\top b)^2.$$

Note for later reference that M_{ab} can also be written as

$$M_{ab} = 1 - \langle 2 \operatorname{diag}(a) - aa^\top, bb^\top \rangle, \tag{3}$$

where $\langle \cdot, \cdot \rangle$ denotes the Frobenius inner product⁵ and $\text{diag}(a)$ is the $n \times n$ diagonal matrix with the entries of a on its diagonal. Let us verify this identity, using $a, b \in \{0, 1\}^n$:

$$\begin{aligned} & 1 - \langle 2 \text{diag}(a) - aa^\top, bb^\top \rangle \\ &= 1 - 2 \langle \text{diag}(a), bb^\top \rangle + \langle aa^\top, bb^\top \rangle \\ &= 1 - 2a^\top b + (a^\top b)^2 = (1 - a^\top b)^2. \end{aligned}$$

Let $\text{suppmat}(M)$ be the binary support matrix of M , so

$$\text{suppmat}(M)_{ab} = \begin{cases} 1 & \text{if } M_{ab} \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

For a given matrix, a *rectangle* is the Cartesian product of a set of row indices and a set of column indices. In [23] it was shown that an exponential number of (monochromatic) rectangles are needed to cover all the 1-entries of the support matrix of M . Equivalently, the corresponding function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ has nondeterministic communication complexity of $\Omega(n)$ bits. For the sake of completeness we repeat the proof here:

Theorem 1 ([23]). *Every 1-monochromatic rectangle cover of $\text{suppmat}(M)$ has size $2^{\Omega(n)}$.*

Proof. Let R_1, \dots, R_k be a 1-cover for f , i.e., a set of (possibly overlapping) 1-monochromatic rectangles in the matrix $\text{suppmat}(M)$ that together cover all 1-entries in $\text{suppmat}(M)$.

We use the following result from [47, Example 3.22 and Section 4.6], which is essentially due to Razborov [60]:

There exist sets $A, B \subseteq \{0, 1\}^n \times \{0, 1\}^n$ and probability distribution μ on $\{0, 1\}^n \times \{0, 1\}^n$ such that all $(a, b) \in A$ have $a^\top b = 0$, all $(a, b) \in B$ have $a^\top b = 1$, $\mu(A) = 3/4$, and there are constants $\alpha, \delta > 0$ (independent of n) such that for all rectangles R ,

$$\mu(R \cap B) \geq \alpha \cdot \mu(R \cap A) - 2^{-\delta n}.$$

(For sufficiently large n , $\alpha = 1/135$ and $\delta = 0.017$ will do.)

Since the R_i are 1-rectangles, they cannot contain elements from B . Hence $\mu(R_i \cap B) = 0$ and $\mu(R_i \cap A) \leq 2^{-\delta n}/\alpha$. However, since all elements of A are covered by the R_i , we have

$$\frac{3}{4} = \mu(A) = \mu\left(\bigcup_{i=1}^k (R_i \cap A)\right) \leq \sum_{i=1}^k \mu(R_i \cap A) \leq k \cdot \frac{2^{-\delta n}}{\alpha}.$$

Hence $k \geq 2^{\Omega(n)}$. □

3 Strong Lower Bounds on Extension Complexity

Here we use the matrix $M = M(n)$ defined in the previous section to prove that the (linear) extension complexity of the cut polytope of the n -vertex complete graph is $2^{\Omega(n)}$, i.e., every (linear) EF of this polytope has an exponential number of inequalities. Then, via reductions, we prove super-polynomial lower bounds for the stable set polytopes and the TSP polytopes. To start, let us define more precisely the slack matrix of a polytope. For a matrix A , let A_i denote the i th row of A and let A^j denote the j th column of A .

Let $P = \{x \in \mathbb{R}^d \mid Ax \leq b\} = \text{conv}(V)$ be a polytope, with $A \in \mathbb{R}^{m \times d}$, $b \in \mathbb{R}^m$ and $V = \{v_1, \dots, v_n\} \subseteq \mathbb{R}^d$. Then $S \in \mathbb{R}_+^{m \times n}$ defined as $S_{ij} := b_i - A_i v_j$ with $i \in [m] := \{1, \dots, m\}$ and $j \in [n] := \{1, \dots, n\}$ is the *slack matrix* of P w.r.t. $Ax \leq b$ and V . We sometimes refer to the submatrix of the slack matrix induced by rows corresponding to facets and columns corresponding to vertices simply as *the* slack matrix of P , denoted by $S(P)$.

Recall that:

⁵The *Frobenius inner product* is the component-wise inner product of two matrices. For matrices X and Y of the same dimensions, this equals $\text{Tr}[X^\top Y]$. When X is symmetric this can also be written $\text{Tr}[XY]$.

1. an *extended formulation* (EF) of P is a linear system in variables (x, y) such that $x \in P$ if and only if there exists y satisfying the system;
2. an *extension* of P is a polytope $Q \subseteq \mathbb{R}^e$ such that there is a linear map $\pi : \mathbb{R}^e \rightarrow \mathbb{R}^d$ with $\pi(Q) = P$;
3. the *extension complexity* of P is the minimum size (i.e., number of inequalities) of an EF of P .

We denote the extension complexity of P by $\text{xc}(P)$.

3.1 The Factorization Theorem

A *rank- r nonnegative factorization* of a (nonnegative) matrix M is a factorization $M = TU$ where T and U are nonnegative matrices with r columns (in case of T) and r rows (in case of U), respectively. The nonnegative rank of M , denoted by $\text{rank}_+(M)$, is thus simply the minimum rank among all nonnegative factorizations of M . Note that $\text{rank}_+(M)$ is also the minimum r such that M is the sum of r nonnegative rank-1 matrices. In particular, the nonnegative rank of a matrix M is at least the nonnegative rank of any submatrix of M .

The following factorization theorem was proved by Yannakakis (see also [29]). It can be stated succinctly as: $\text{xc}(P) = \text{rank}_+(S)$ whenever P is a polytope and S a slack matrix of P . We include a sketch of the proof for completeness and we will use the following lemma which follows easily from Farkas's Lemma [64, 73] by first showing that $\mathbf{0}^\top x \leq 1$ can be derived from the system.

Lemma 2. *Let $P = \{x \in \mathbb{R}^d \mid Ax \leq b\}$ be a (possibly unbounded) polyhedron that admits a direction $u \in \mathbb{R}^d$ with $-\infty < \min\{u^\top x \mid x \in P\} < \max\{u^\top x \mid x \in P\} < +\infty$, and $c^\top x \leq \delta$ a valid inequality for P . Then there exist nonnegative multipliers $\lambda \in \mathbb{R}^d$ such that $\lambda^\top A = c^\top$ and $\lambda^\top b = \delta$, i.e., $c^\top x \leq \delta$ can be derived as a nonnegative combination from $Ax \leq b$. In particular, this holds whenever P is a polytope of dimension at least 1 or whenever P is an unbounded polyhedron that linearly projects to a polytope of dimension at least 1.*

We are ready to state Yannakakis's factorization theorem.

Theorem 3 ([72]). *Let $P = \{x \in \mathbb{R}^d \mid Ax \leq b\} = \text{conv}(V)$ be a polytope with $\dim(P) \geq 1$, and let S denote the slack matrix of P w.r.t. $Ax \leq b$ and V . Then the following are equivalent for all positive integers r :*

- (i) S has nonnegative rank at most r ;
- (ii) P has an extension of size at most r (that is, with at most r facets);
- (iii) P has an EF of size at most r (that is, with at most r inequalities).

(*sketch*). It should be clear that (ii) implies (iii). We prove that (i) implies (ii), and then that (iii) implies (i).

First, consider a rank- r nonnegative factorization $S = TU$ of the slack matrix of P . Notice that we may assume that no column of T is zero, because otherwise r can be decreased. We claim that P is the image of

$$Q := \{(x, y) \in \mathbb{R}^{d+r} \mid Ax + Ty = b, y \geq \mathbf{0}\}$$

under the projection π_x onto the x -space. We see immediately that $\pi_x(Q) \subseteq P$ since $Ty \geq \mathbf{0}$. To prove the inclusion $P \subseteq \pi_x(Q)$, it suffices to remark that for each point $v_j \in V$ the point (v_j, U^j) is in Q since

$$Av_j + TU^j = Av_j + b - Av_j = b \text{ and } U^j \geq \mathbf{0}.$$

Since no column of T is zero, Q is a polytope. Moreover, Q has at most r facets, and is thus an extension of P of size at most r . This proves that (i) implies (ii).

Second, suppose that the system

$$E^-x + F^-y = g^-, \quad E^\leq x + F^\leq y \leq g^\leq$$

with $(x, y) \in \mathbb{R}^{d+r}$ defines an EF of P . Let $Q \subseteq \mathbb{R}^{d+r}$ denote the set of solutions to this system. Thus Q is a (not necessarily bounded) polyhedron. For each point $v_j \in V$, pick $w_j \in \mathbb{R}^r$ such that $(v_j, w_j) \in Q$. Because

$$Ax \leq b \iff \exists y : E^=x + F^=y = g^=, E^{\leq}x + F^{\leq}y \leq g^{\leq},$$

each inequality in $Ax \leq b$ is valid for all points of Q . Let S_Q be the nonnegative matrix that records the slacks of the points (v_j, w_j) with respect to the inequalities of $E^{\leq}x + F^{\leq}y \leq g^{\leq}$, and then of $Ax \leq b$. By construction, the submatrix obtained from S_Q by deleting the r first rows is exactly S , thus $\text{rank}_+(S) \leq \text{rank}_+(S_Q)$. Furthermore, by Lemma 2 any valid inequality $c^T x \leq \delta$ is a nonnegative combination of inequalities of the system $Ax \leq b$ and thus every row of S_Q is a nonnegative combination of the first r rows of S_Q . Thus, $\text{rank}_+(S_Q) \leq r$. Therefore, $\text{rank}_+(S) \leq r$. Hence (iii) implies (i). \square

Remark 1. By the factorization theorem, if polytope $P \subseteq \mathbb{R}^d$ has an EF of size r , then its slack matrix S has a nonnegative factorization $S = TU$ of rank r . But then $Ax + Ty = b$, $y \geq \mathbf{0}$ is an EF of P in slack form with $d+r$ variables, r inequalities and m equalities, where m is the number of rows in the linear description $Ax \leq b$ of P . Notice that if $m > d+r$ some of these equalities will be redundant, and that there always exists a subset of at most $d+r$ equalities defining the same subspace. By removing redundant equalities from the EF, we can assume that there are at most $d+r$ equalities in the EF.

We will prove a generalization of Theorem 3 for arbitrary closed convex cones in Section 4, but for now this special case is all we need.

We would like to emphasize that we will not restrict the slack matrix to have rows corresponding only to the facet-defining inequalities. This is not an issue since appending rows corresponding to redundant⁶ inequalities does not change the nonnegative rank of the slack matrix. This fact was already used in the second part of the previous proof (see also Lemma 15 below).

Theorem 3 shows in particular that we can lower bound the extension complexity of P by lower bounding the nonnegative rank of its slack matrix S ; in fact it suffices to lower bound the nonnegative rank of any submatrix of the slack matrix S corresponding to an implied system of inequalities. To that end, Yannakakis made the following connection with nondeterministic communication complexity. Again, we include the (easy) proof for completeness.

Theorem 4 ([72]). *Let M be any matrix with nonnegative real entries and $\text{suppmat}(M)$ its support matrix. Then $\text{rank}_+(M)$ is lower bounded by the rectangle covering bound for $\text{suppmat}(M)$.*

Proof. If $M = TU$ is a rank- r nonnegative factorization of M , then S can be written as the sum of r nonnegative rank-1 matrices:

$$S = \sum_{k=1}^r T^k U_k.$$

Taking the support on each side, we find

$$\begin{aligned} \text{supp}(S) &= \bigcup_{k=1}^r \text{supp}(T^k U_k) \\ &= \bigcup_{k=1}^r \text{supp}(T^k) \times \text{supp}(U_k). \end{aligned}$$

Therefore, $\text{suppmat}(M)$ has a 1-monochromatic rectangle cover with r rectangles. \square

⁶An inequality of a linear system is called *redundant* if removing the inequality from the system does not change the set of solutions.

3.2 Cut and Correlation Polytopes

Let $K_n = (V_n, E_n)$ denote the n -vertex complete graph. For a set X of vertices of K_n , we let $\delta(X)$ denote the set of edges of K_n with one endpoint in X and the other in its complement \bar{X} . This set $\delta(X)$ is known as the *cut* defined by X . For a subset F of edges of K_n , we let $\chi^F \in \mathbb{R}^{E_n}$ denote the *characteristic vector* of F , with $\chi_e^F = 1$ if $e \in F$ and $\chi_e^F = 0$ otherwise. The *cut polytope* $\text{CUT}(n)$ is defined as the convex hull of the characteristic vectors of all cuts in the complete graph $K_n = (V_n, E_n)$. That is,

$$\text{CUT}(n) := \text{conv}\{\chi^{\delta(X)} \in \mathbb{R}^{E_n} \mid X \subseteq V_n\}.$$

We will not deal with the cut polytopes directly, but rather with 0/1-polytopes that are linearly isomorphic to them. The *correlation polytope* (or *Boolean quadric polytope*) $\text{COR}(n)$ is defined as the convex hull of all the rank-1 binary symmetric matrices of size $n \times n$. In other words,

$$\text{COR}(n) := \text{conv}\{bb^\top \in \mathbb{R}^{n \times n} \mid b \in \{0, 1\}^n\}.$$

We use the following known result:

Theorem 5 ([21]). *For all n , $\text{COR}(n)$ is linearly isomorphic to $\text{CUT}(n+1)$.*

Because M is nonnegative, Eq. (3) gives us a linear inequality that is satisfied by all vertices bb^\top of $\text{COR}(n)$, and hence (by convexity) is satisfied by all points of $\text{COR}(n)$:

Lemma 6. *For all $a \in \{0, 1\}^n$, the inequality*

$$\langle 2 \text{diag}(a) - aa^\top, x \rangle \leq 1 \tag{4}$$

is valid for $\text{COR}(n)$. Moreover, the slack of vertex $x = bb^\top$ with respect to (4) is precisely M_{ab} .

We remark that (4) is weaker than the *hypermetric inequality* [24] $\langle \text{diag}(a) - aa^\top, x \rangle \leq 0$, in the sense that the face defined by the former is strictly contained in the face defined by the latter. Nevertheless, we persist in using (4). Now, we prove the main result of this section.

Theorem 7. *There exists some constant $C > 0$ such that, for all n ,*

$$\text{xc}(\text{CUT}(n+1)) = \text{xc}(\text{COR}(n)) \geq 2^{Cn}.$$

In particular, the extension complexity of $\text{CUT}(n)$ is $2^{\Omega(n)}$.

Proof. The equality is implied by Theorem 5. Now, consider any system of linear inequalities describing $\text{COR}(n)$ starting with the 2^n inequalities (4), and a slack matrix S w.r.t. this system and $\{bb^\top \mid b \in \{0, 1\}^n\}$. Next delete from this slack matrix all rows except the 2^n first rows. By Lemma 6, the resulting $2^n \times 2^n$ matrix is M . Using Theorems 3, 4, and 1, and the fact that the nonnegative rank of a matrix is at least the nonnegative rank of any of its submatrices, we have

$$\begin{aligned} \text{xc}(\text{COR}(n)) &= \text{rank}_+(S) \\ &\geq \text{rank}_+(M) \\ &\geq 2^{Cn} \end{aligned}$$

for some positive constant C . □

3.3 Stable Set Polytopes

A *stable set* S (also called an *independent set*) of a graph $G = (V, E)$ is a subset $S \subseteq V$ of the vertices such that no two of them are adjacent. For a subset $S \subseteq V$, we let $\chi^S \in \mathbb{R}^V$ denote the *characteristic vector* of S , with $\chi_v^S = 1$ if $v \in S$ and $\chi_v^S = 0$ otherwise. The *stable set polytope*, denoted $\text{STAB}(G)$, is the convex hull of the characteristic vectors of all stable sets in G , i.e.,

$$\text{STAB}(G) := \text{conv}\{\chi^S \in \mathbb{R}^V \mid S \text{ stable set of } G\}.$$

Recall that a polytope Q is an extension of a polytope P if P is the image of Q under a linear projection.

Lemma 8. For each n , there exists a graph H_n with $O(n^2)$ vertices such that $\text{STAB}(H_n)$ contains a face that is an extension of $\text{COR}(n)$.

Proof. Consider the complete graph K_n with vertex set $V_n := [n]$. For each vertex i of K_n we create two vertices labeled ii, \bar{ii} in H_n and an edge between them. For each edge ij of K_n , we add to H_n four vertices labeled $ij, \bar{ij}, \underline{ij}, \underline{\bar{ij}}$ and all possible six edges between them. We further add the following eight edges to H_n :

$$\begin{aligned} & \{ij, \bar{ii}\}, \{ij, \bar{jj}\}, \{\bar{ij}, ii\}, \{\bar{ij}, \bar{jj}\}, \\ & \{\underline{ij}, \bar{ii}\}, \{\underline{ij}, jj\}, \{\underline{\bar{ij}}, ii\}, \{\underline{\bar{ij}}, jj\}. \end{aligned}$$

See Fig. 1 for an illustration. The number of vertices in H_n is $2n + 4\binom{n}{2}$.

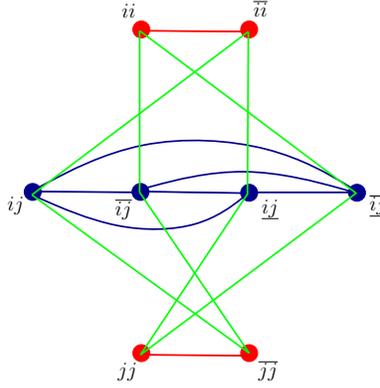


Figure 1: The edges and vertices of H_n corresponding to vertices i, j and edge ij of K_n .

Thus the vertices and edges of K_n are represented by cliques of size 2 and 4 respectively in H_n . We will refer to these as *vertex-cliques* and *edge-cliques* respectively. Consider the face $F = F(n)$ of $\text{STAB}(H_n)$ whose vertices correspond to the stable sets containing exactly one vertex in each vertex-clique and each edge-clique. (The vertices in this face correspond exactly to stable sets of H_n with maximum cardinality.)

Consider the linear map $\pi : \mathbb{R}^{V(H_n)} \rightarrow \mathbb{R}^{n \times n}$ mapping a point $x \in \mathbb{R}^{V(H_n)}$ to the point $y \in \mathbb{R}^{n \times n}$ such that $y_{ij} = y_{ji} = x_{ij}$ for $i \leq j$. In this equation, the subscripts in y_{ij} and y_{ji} refer to an ordered pair of elements in $[n]$, while the subscript in x_{ij} refers to a vertex of H_n that corresponds either to a vertex of K_n (if $i = j$) or to an edge of K_n (if $i \neq j$).

We claim that the image of F under π is $\text{COR}(n)$, hence F is an extension of $\text{COR}(n)$; observe that it suffices to consider 0/1 vertices as F is a 0/1 polytope and the projection is an orthogonal projection. Indeed, pick an arbitrary stable set S of H_n such that $x := \chi^S$ is on face F . Then define $b \in \{0, 1\}^n$ by letting $b_i := 1$ if $ii \in S$ and $b_i := 0$ otherwise (i.e., $\bar{ii} \in S$). Notice that for the edge ij of K_n we have $ij \in S$ if and only if both vertices ii and jj belong to S . Hence, $\pi(x) = y = bb^\top$ is a vertex of $\text{COR}(n)$. This proves $\pi(F) \subseteq \text{COR}(n)$. Now pick a vertex $y := bb^\top$ of $\text{COR}(n)$ and consider the unique maximum stable set S that contains vertex ii if $b_i = 1$ and vertex \bar{ii} if $b_i = 0$. Then $x := \chi^S$ is a vertex of F with $\pi(x) = y$. Hence, $\pi(F) \supseteq \text{COR}(n)$. Thus $\pi(F) = \text{COR}(n)$. This concludes the proof. \square

Our next lemma establishes simple monotonicity properties of the extension complexity used in our reduction.

Lemma 9. Let P, Q and F be polytopes. Then the following hold:

- (i) if F is an extension of P , then $\text{xc}(F) \geq \text{xc}(P)$;
- (ii) if F is a face of Q , then $\text{xc}(Q) \geq \text{xc}(F)$.

Proof. The first part is obvious because every extension of F is in particular an extension of P . For the second part, notice that a slack matrix of F can be obtained from the (facet-vs-vertex) slack matrix of Q by deleting columns corresponding to vertices not in F . Now apply Theorem 3. \square

Using previous results, we can prove the following result about the worst-case extension complexity of the stable set polytope.

Theorem 10. *For all n , one can construct a graph G_n with n vertices such that the extension complexity of the stable set polytope $\text{STAB}(G_n)$ is $2^{\Omega(n^{1/2})}$.*

Proof. W.l.o.g., we may assume $n \geq 18$. For an integer $p \geq 3$, let $f(p) := |V(H_p)| = 2p + 4\binom{p}{2}$. Given $n \geq 18$, we define p as the largest integer with $f(p) \leq n$. Now let G_n be obtained from H_p by adding $n - f(p)$ isolated vertices. Then $\text{STAB}(H_p)$ is linearly isomorphic to a face of $\text{STAB}(G_n)$. Using Theorem 7 in combination with Lemmas 8 and 9, we find that

$$\begin{aligned} \text{xc}(\text{STAB}(G_n)) &\geq \text{xc}(\text{STAB}(H_p)) \\ &\geq \text{xc}(\text{COR}(p)) \\ &= 2^{\Omega(p)} \\ &= 2^{\Omega(n^{1/2})}. \end{aligned}$$

\square

3.4 TSP Polytopes

Recall that $\text{TSP}(n)$, the *traveling salesman polytope* or *TSP polytope* of $K_n = (V_n, E_n)$, is defined as the convex hull of the characteristic vectors of all subsets $F \subseteq E_n$ that define a tour of K_n . That is,

$$\text{TSP}(n) := \text{conv}\{\chi^F \in \mathbb{R}^{E_n} \mid F \subseteq E_n \text{ is a tour of } K_n\}.$$

We now prove that the polytope $\text{COR}(n)$ is the linear projection of a face of $\text{TSP}(O(n^2))$, implying the following:

Lemma 11. *For each n , there exists a positive integer $q = O(n^2)$ vertices such that $\text{TSP}(q)$ contains a face that is an extension of $\text{COR}(n)$.*

Proof. Recall that

$$\text{COR}(n) = \text{conv}\{bb^\top \in \mathbb{R}^{n \times n} \mid b \in \{0, 1\}^n\}.$$

To prove the lemma we start with constructing a graph G_n with $q = O(n^2)$ vertices such that the tours of G_n correspond to the $n \times n$ rank-1 binary symmetric matrices bb^\top , where $b \in \{0, 1\}^n$. This is done in three steps:

- (i) define a 3SAT formula ϕ_n with n^2 variables such that the satisfying assignments of ϕ_n bijectively correspond to the matrices bb^\top , where $b \in \{0, 1\}^n$;
- (ii) construct a directed graph D_n with $O(n^2)$ vertices such that each directed tour of D_n defines a satisfying assignment of ϕ_n , and conversely each satisfying assignment of ϕ_n has at least one corresponding directed tour in D_n ;
- (iii) modify the directed graph D_n into an undirected graph G_n in such a way that the tours of G_n bijectively correspond to the directed tours of D_n .

Step (i). For defining ϕ_n we use Boolean variables $C_{ij} \in \{0, 1\}$ for $i, j \in [n]$ and let

$$\phi_n := \bigwedge_{\substack{i, j \in [n] \\ i \neq j}} [(C_{ii} \vee C_{jj} \vee \overline{C_{ij}}) \wedge (C_{ii} \vee \overline{C_{jj}} \vee \overline{C_{ij}}) \wedge (\overline{C_{ii}} \vee C_{jj} \vee \overline{C_{ij}}) \wedge (\overline{C_{ii}} \vee \overline{C_{jj}} \vee C_{ij})].$$

The four clauses $(C_{ii} \vee C_{jj} \vee \overline{C_{ij}})$, $(C_{ii} \vee \overline{C_{jj}} \vee \overline{C_{ij}})$, $(\overline{C_{ii}} \vee C_{jj} \vee \overline{C_{ij}})$ and $(\overline{C_{ii}} \vee \overline{C_{jj}} \vee C_{ij})$ model the equation $C_{ij} = C_{ii} \wedge C_{jj}$. Hence, $C \in \{0, 1\}^{n \times n}$ satisfies ϕ_n if and only if there exists $b \in \{0, 1\}^n$ such that $C_{ij} = b_i \wedge b_j$ for all $i, j \in [n]$, or in matrix language, $C = bb^\top$.

Step (ii). To construct a directed graph D_n whose directed tours correspond to the satisfying assignments of ϕ_n we use the standard reduction from 3SAT to HAMPATH [67].

We order the variables of ϕ_n arbitrarily and construct a gadget for each variable as follows. Suppose that the k th variable occurs in p clauses. We create a chain of $3p + 1$ nodes, labeled $v_{k,1}, \dots, v_{k,3p+1}$, where each node $v_{k,\ell}$ with $\ell < 3p + 1$ is connected to the next node $v_{k,\ell+1}$ with two opposite directed edges. Figure 2 illustrates this. Traversing this chain from left to right is interpreted as setting the k th variable to false and traversing it from right to left is interpreted as setting the k th variable to true. We also have two nodes s_k, t_k connected to this chain with directed edges $(s_k, v_{k,1})$, $(s_k, v_{k,3p+1})$, $(v_{k,1}, t_k)$ and $(v_{k,3p+1}, t_k)$ creating a diamond structure.

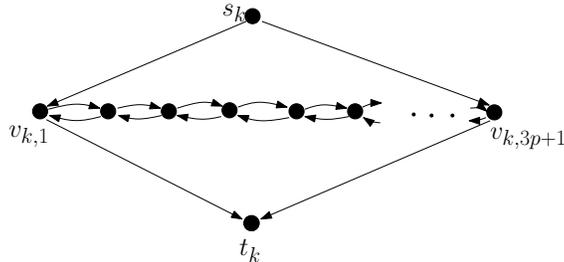


Figure 2: Gadget for the k th variable occurring in p clauses.

Next, we order the clauses of ϕ_n arbitrarily and create a node for each clause. The node for the m th clause is denoted by w_m . We connect these extra nodes to the gadgets for the variables as follows. Suppose, as before, that the k th variable appears in p clauses. Consider the ℓ th of these clauses in which the k th variable appears, and let m be the index of that clause. If the k th variable appears negated in the m th clause then we add the path $v_{k,3\ell-1}, w_m, v_{k,3\ell}$. Otherwise the k th variable appears non-negated in the m th clause and we add the path $v_{k,3\ell}, w_m, v_{k,3\ell-1}$. Figure 3 illustrates this.

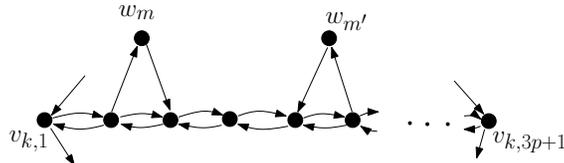


Figure 3: Gadgets for clauses in case the k th variable appears negated in the m th clause and non-negated in the m' th clause.

Next we connect the gadgets corresponding to the variables by identifying t_k with s_{k+1} for $1 \leq k < n^2$. Finally, we add a directed edge from t_{n^2} to s_1 . Figure 4 illustrates the final directed graph obtained.

To see why the directed tours of the final directed graph D_n define satisfying assignments of our Boolean formula ϕ_n , observe that each directed tour of D_n encodes a truth assignment to the n^2 variables depending on which way the corresponding chains are traversed. Because a directed tour visits every node and because the node w_m corresponding to a clause can be visited only if we satisfy it, the truth assignment satisfies ϕ_n . Conversely, every satisfying assignment of ϕ_n yields at least one directed tour in D_n . (If the m th clause is satisfied by the value of more than one variable, we visit w_m only once, from the chain of the first variable whose value makes the clause satisfied).

Step (iii). For each node v of D_n we create a path $v_{\text{in}}, v_{\text{mid}}, v_{\text{out}}$ in the (undirected) graph G_n . For each directed edge (v, w) of D_n , we add to the graph G_n an edge between v_{out} and w_{in} . As is easily seen, the

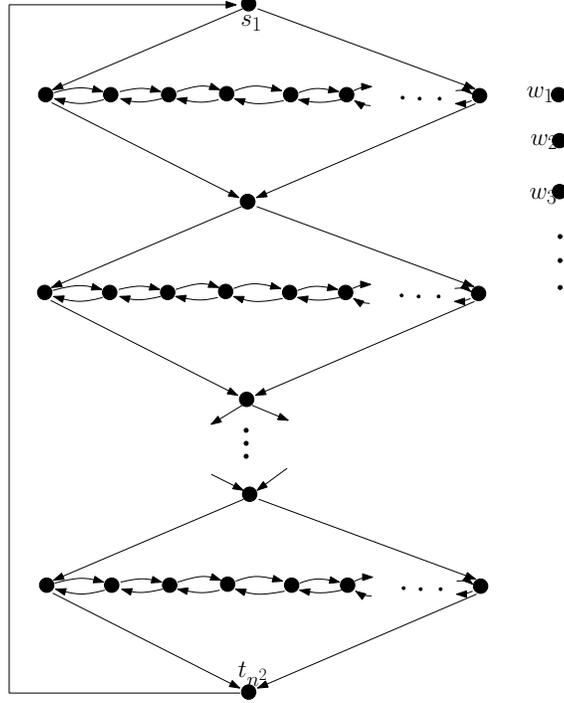


Figure 4: Final graph. Directed edges incident to nodes w_m depend on the actual ordering of variables and clauses in Boolean formula ϕ_n .

tours of G_n bijectively correspond to the directed tours of D_n . Note that G_n has $q := 3(n(n-1) \cdot 13 + n \cdot (3n-2) + n^2 + 4n(n-1)) = O(n^2)$ vertices.

Consider the face F of $\text{TSP}(q)$ defined by setting to 0 all variables x_e corresponding to non-edges of G_n , so that the vertices of F are the characteristic vectors of the tours of G_n . To conclude the proof we give a linear projection $\pi : x \mapsto y := \pi(x)$ mapping F to $\text{COR}(n)$. For $x \in \mathbb{R}^{E_q}$ and $i, j \in [n]$, we let $y_{ij} = x_e$, where e is the edge $(v_{out,k,2}, v_{in,k,1})$ of G_n corresponding to the directed edge $(v_{k,2}, v_{k,1})$ and k is the index of the variable C_{ij} of ϕ_n . It follows from the above discussion that π maps the face F of $\text{TSP}(q)$ to $\text{COR}(n)$. The lemma follows. \square

The final theorem in this section follows from Theorem 7, Lemmas 9 and 11, using an argument similar to that used in the proof of Theorem 10.

Theorem 12. *The extension complexity of the TSP polytope $\text{TSP}(n)$ is $2^{\Omega(n^{1/2})}$.*

4 Conic and Semidefinite EFs

In this section we extend Yannakakis's factorization theorem (Theorem 3) to arbitrary closed convex cones. Recall that every EF can be brought into slack form $Ex + Fy = g, y \geq \mathbf{0}$ without increasing its size (see Remark 1). This form is the basis of our generalization: we replace the nonnegativity constraint $y \geq \mathbf{0}$ by a general conic constraint $y \in C$.

Let $Q = \{(x, y) \in \mathbb{R}^{d+k} \mid Ex + Fy = g, y \in C\}$ for an arbitrary closed convex cone $C \subseteq \mathbb{R}^k$, where $E \in \mathbb{R}^{p \times d}, F \in \mathbb{R}^{p \times k}$, and $g \in \mathbb{R}^p$. Let $C^* := \{z \in \mathbb{R}^k \mid z^\top y \geq 0, \forall y \in C\}$ denote the *dual cone* of C . We define the *projection cone* of Q as

$$C_Q := \{\mu \in \mathbb{R}^p \mid F^\top \mu \in C^*\}$$

and

$$\text{proj}_x(Q) := \{x \in \mathbb{R}^d \mid \mu^\top Ex \leq \mu^\top g, \forall \mu \in C_Q\}.$$

In a first step we show that $\text{proj}_x(Q)$ equals

$$\pi_x(Q) := \{x \in \mathbb{R}^d \mid \exists y \in \mathbb{R}^k : (x, y) \in Q\},$$

the projection of Q to the x -space.

Lemma 13. *With the above notation, we have $\pi_x(Q) = \text{proj}_x(Q)$.*

Proof. Let $\alpha \in \pi_x(Q)$. Then there exists $y \in C$ with $E\alpha + Fy = g$. Pick any $\mu \in C_Q$. Then, $\mu^\top E\alpha + \mu^\top Fy = \mu^\top g$ holds. Since $F^\top \mu \in C^*$ and $y \in C$ we have that $(F^\top \mu)^\top y = \mu^\top Fy \geq 0$. Therefore $\mu^\top E\alpha \leq \mu^\top g$ holds for all $\mu \in C_Q$. We conclude $\alpha \in \text{proj}_x(Q)$ and as α was arbitrary $\pi_x(Q) \subseteq \text{proj}_x(Q)$ follows.

Now suppose $\alpha \notin \pi_x(Q)$. We show that this implies $\alpha \notin \text{proj}_x(Q)$. As $\alpha \notin \pi_x(Q)$ there is no $y \in C$ such that $Fy = g - E\alpha$ or, equivalently, the convex cone $F(C) := \{Fy \mid y \in C\}$ does not contain the point $g - E\alpha$. Since C is a closed cone, so is $F(C)$. Therefore, by the Strong Separation Theorem there exists $\mu \in \mathbb{R}^p$ such that $\mu^\top z \geq 0$ is valid for $F(C)$ but $\mu^\top(g - E\alpha) < 0$. Then $\mu^\top z = \mu^\top(Fy) = (\mu^\top F)y \geq 0$ is valid for C , i.e., $(\mu^\top F)y \geq 0$ holds for all $y \in C$, implying $F^\top \mu \in C^*$ and so $\mu \in C_Q$. Because $\mu^\top(g - E\alpha) < 0$ we have $\mu^\top E\alpha > \mu^\top g$ with $\mu \in C_Q$ so that $\alpha \notin \text{proj}_x(Q)$ follows. \square

Let $P = \{x \in \mathbb{R}^d \mid Ax \leq b\} = \text{conv}(V)$ be a polytope, with $A \in \mathbb{R}^{m \times d}$, $b \in \mathbb{R}^m$ and $V = \{v_1, \dots, v_n\} \subseteq \mathbb{R}^d$. Then $Ex + Fy = g$, $y \in C$ is a *conic EF* of P whenever $x \in P$ if and only if there exists $y \in C$ such that $Ex + Fy = g$. The set $Q = \{(x, y) \in \mathbb{R}^{d+k} \mid Ex + Fy = g, y \in C\}$ is then called a *conic extension* of P w.r.t. C .

We now prove a factorization theorem for the slack matrix of polytopes that generalizes Yannakakis's factorization theorem in the linear case (Theorem 3). Yannakakis's result can be obtained as a corollary of our result by taking $C = \mathbb{R}_+^k$, and using Theorem 14 together with the fact that $(\mathbb{R}_+^k)^* = \mathbb{R}_+^k$.

Theorem 14. *Let $P = \{x \in \mathbb{R}^d \mid Ax \leq b\} = \text{conv}(V)$ be a polytope with $\dim(P) \geq 1$ defined by m inequalities and n points respectively, and let S be the slack matrix of P w.r.t. $Ax \leq b$ and V . Also, let $C \subseteq \mathbb{R}^k$ be a closed convex cone. Then, the following are equivalent:*

(i) *There exist T, U such that (the transpose of) each row of T is in C^* , each column of U is in C , and $S = TU$.*

(ii) *There exists a conic extension $Q = \{(x, y) \in \mathbb{R}^{d+k} \mid Ex + Fy = g, y \in C\}$ with $P = \pi_x(Q)$.*

Before proving the theorem, we prove a lemma which will allow us to get rid of rows of a slack matrix that correspond to redundant inequalities. Below, we call a factorization as in (i) a factorization of S w.r.t. C .

Lemma 15. *Let $P \subseteq \mathbb{R}^d$ be a polytope with $\dim(P) \geq 1$, let S and S' be two slack matrices of P , and let $C \subseteq \mathbb{R}^k$ be a closed convex cone. Then S has a factorization w.r.t. C if and only if S' has a factorization w.r.t. C .*

Proof. It suffices to prove the theorem when S' is the submatrix of S induced by the rows corresponding to facet-defining inequalities and the columns corresponding to vertices, that is, when $S' = S(P)$. One implication is clear: if S has a factorization w.r.t. C , then S' also because S' is a submatrix of S .

For the other implication, consider a system $Ax \leq b$ of m inequalities and a set $V = \{v_1, \dots, v_n\}$ of n points such that $P = \{x \in \mathbb{R}^d \mid Ax \leq b\} = \text{conv}(V)$. Assume that the f first inequalities of $Ax \leq b$ are facet-defining, while the remaining $m - f$ are not, and that the v first points of V are vertices, while the remaining $n - v$ are not.

Consider an inequality $A_i x \leq b_i$ with $i > f$. Suppose first that the inequality is redundant. By Farkas's lemma (using $\dim(P) \geq 1$), there exist nonnegative coefficients $\mu_{i,k}$ with $k \in [f]$ such that $A_i = \sum_{k \in [f]} \mu_{i,k} A_k$ and $b_i = \sum_{k \in [f]} \mu_{i,k} b_k$ as P is a polytope. If the inequality is not redundant, since it is not facet-defining,

it is satisfied with equality by all points of P . In this case, we let $\mu_{i,k} := 0$ for all $k \in [f]$. Finally, for $i \leq f$ we let $\mu_{i,k} := 1$ if $i = k$ and $\mu_{i,k} := 0$ otherwise.

Next, consider a point v_j with $j > v$. Because v_j is in P , it can be expressed as a convex combination of the vertices of P : $v_j = \sum_{\ell \in [v]} \lambda_{j,\ell} v_\ell$, where $\lambda_{j,\ell}$ with $\ell \in [v]$ are nonnegative coefficients that sum up to 1. Similarly as above, for $j \leq v$ we let $\lambda_{j,\ell} := 1$ if $j = \ell$ and $\lambda_{j,\ell} := 0$ otherwise.

Now, let $S' = TU$ be a factorization of S' w.r.t. C . That is, we have row vectors T_1, \dots, T_f with $(T_k)^\top \in C^*$ (for $k \in [f]$) and column vectors U^1, \dots, U^v with $U^\ell \in C$ (for $\ell \in [v]$) such that $b_k - A_k v_\ell = S'_{k\ell} = T_k U^\ell$ for $k \in [f], \ell \in [v]$.

We extend the factorization of S' into a factorization of S by letting $T_i := \sum_{k \in [f]} \mu_{i,k} T_k$ and $U^j := \sum_{\ell \in [v]} \lambda_{j,\ell} U^\ell$ for $i > f$ and $j > v$. Given our choice of coefficients, these equations also hold for $i \leq f$ and $j \leq v$. Clearly, each T_i (transposed) is in C^* and each U^j is in C . A straightforward computation then shows $T_i U^j = S_{ij}$ for all $i \in [m], j \in [n]$. Therefore, T_i with $i \in [m]$ and U^j with $j \in [n]$ define a factorization of S w.r.t. C . \square

of Theorem 14. We first show that a factorization induces a conic extension. Suppose there exist matrices T, U as above. We claim that Q with $E := A, F := T$ and $g := b$ has the desired properties. Let $v_j \in V$, then $S^j = TU^j = b - Av_j$ and so it follows that $(v_j, U^j) \in Q$ and $v_j \in \pi_x(Q)$. Now let $x \in \pi_x(Q)$. Then, there exists $y \in C$ such that $Ax + Ty = b$. Since $T_i y \geq 0$ for all $i \in [m]$, we have that $x \in P$. This proves the first implication.

For the converse, suppose $P = \pi_x(Q)$ with Q being a conic extension of P . By Lemma 13, $\pi_x(Q) = \{x \in \mathbb{R}^d \mid \mu^\top E x \leq \mu^\top g, \forall \mu \in C_Q\}$, where $C_Q = \{\mu \in \mathbb{R}^p \mid F^\top \mu \in C^*\}$. By Lemma 15, it suffices to prove that the submatrix of S induced by the rows corresponding to the inequalities of $Ax \leq b$ that define facets of P admits a factorization w.r.t. C . Thus, we assume for the rest of the proof that all rows of S correspond to facets of P . Then, for any facet-defining inequality $A_i x \leq b_i$ of P there exists $\mu_i \in C_Q$ such that $\mu_i^\top E x \leq \mu_i^\top g$ defines the same facet as $A_i x \leq b_i$. (This follows from the fact that C_Q is closed; see also [51, Theorem 4.3.4].) Scaling μ_i if necessary, this means that $\mu_i^\top E = A_i + c^\top$ and $\mu_i^\top g = b_i + \delta$, where $c^\top x = \delta$ is satisfied for all points of P . We define $T_i := \mu_i^\top F$ for all i ; in particular $(T_i)^\top \in C^*$ as $\mu_i \in C_Q$. Now let $v_j \in V$. Since $P = \pi_x(Q)$, there exists a $y_j \in C$ such that $E v_j + F y_j = g$ and so $\mu_i^\top E v_j + \mu_i^\top F y_j = \mu_i^\top g$. With the above we have $A_i v_j + c^\top v_j + T_i y_j = b_i + \delta$, hence $A_i v_j + T_i y_j = b_i$ and as $v_j \in \pi_x(Q)$ we deduce $T_i y_j \geq 0$. The slack of v_j w.r.t. $A_i x \leq b_i$ is $b_i - A_i v_j = \mu_i^\top g - \mu_i^\top E v_j = \mu_i^\top F y_j = T_i y_j$. This implies the factorization $S = TU$ with $T_i = \mu_i^\top F$ and $U^j = y_j$. \square

For a positive integer r , we let \mathbb{S}_+^r denote the cone of $r \times r$ symmetric positive semidefinite matrices embedded in $\mathbb{R}^{r(r+1)/2}$ in such a way that, for all $y, z \in \mathbb{S}_+^r$, the scalar product $z^\top y$ is the Frobenius product of the corresponding matrices. A *semidefinite EF (resp. extension)* of size r is simply a conic EF (resp. extension) w.r.t. $C = \mathbb{S}_+^r$. The *semidefinite extension complexity* of polytope P , denoted by $\text{xc}_{SDP}(P)$, is the minimum r such that P has a semidefinite EF of size r . Observe that $(\mathbb{S}_+^r)^* = \mathbb{S}_+^r$. Hence, taking $C := \mathbb{S}_+^k$ and $k := r(r+1)/2$ in Theorem 14, we obtain the following factorization theorem for semidefinite EFs.

Corollary 16. *Let $P = \{x \in \mathbb{R}^d \mid Ax \leq b\} = \text{conv}(V)$ be a polytope. Then the slack matrix S of P w.r.t. $Ax \leq b$ and V has a factorization $S = TU$ so that $(T_i)^\top, U^j \in \mathbb{S}_+^r$ if and only if there exists a semidefinite extension $Q = \{(x, y) \in \mathbb{R}^{d+r(r+1)/2} \mid Ex + Fy = g, y \in \mathbb{S}_+^r\}$ with $P = \pi_x(Q)$.*

Analogous to nonnegative factorizations and nonnegative rank, we can define PSD factorizations and PSD rank. A *rank- r PSD factorization* of an $m \times n$ matrix M is a collection of $r \times r$ symmetric positive semidefinite matrices T_1, \dots, T_m and U^1, \dots, U^n such that the Frobenius product $\langle T_i, U^j \rangle = \text{Tr} [(T_i)^\top U^j] = \text{Tr} [T_i U^j]$ equals M_{ij} for all $i \in [m], j \in [n]$. The *PSD rank* of M is the minimum r such that M has a rank- r PSD factorization. We denote this by $\text{rank}_{PSD}(M)$. By Corollary 16 (and also Lemma 15), the semidefinite extension complexity of a polytope P is equal to the PSD rank of any slack matrix of P : $\text{xc}_{SDP}(P) = \text{rank}_{PSD}(S)$ whenever S is a slack matrix of P . In the next section we will show that $\text{rank}_{PSD}(M)$ can be expressed in terms of the amount of communication needed by a one-way quantum communication protocol for computing M in expectation (Corollary 18).

5 Quantum Communication and PSD Factorizations

In this section we explain the connection with quantum communication. This yields results that are interesting in their own right, and also clarifies where the matrix M of Section 2 came from.

For a general introduction to quantum computation we refer to [57] and to [56], and for quantum communication complexity we refer to [22] and to [15]. For our purposes, an r -dimensional *quantum state* ρ is an $r \times r$ PSD matrix of trace 1.⁷ A k -qubit state is a state in dimension $r = 2^k$. If ρ has rank 1, it can be written as an outer product $|\phi\rangle\langle\phi|$ for some unit vector $|\phi\rangle$, which is sometimes called a *pure state*. We use $|i\rangle$ to denote the pure state vector that has 1 at position i and 0s elsewhere. A quantum measurement (POVM) is described by a set of PSD matrices $\{E_\theta\}_{\theta \in \Theta}$, each labeled by a real number θ , and summing to the r -dimensional identity: $\sum_{\theta \in \Theta} E_\theta = I$. When measuring state ρ with this measurement, the probability of outcome θ equals $\text{Tr}[E_\theta \rho]$. Note that if we define the PSD matrix $E := \sum_{\theta \in \Theta} \theta E_\theta$, then the *expected value* of the measurement outcome is $\sum_{\theta \in \Theta} \theta \text{Tr}[E_\theta \rho] = \text{Tr}[E \rho]$.

5.1 Quantum Protocols

A *one-way quantum protocol with r -dimensional messages* can be described as follows. On input i , Alice sends Bob an r -dimensional state ρ_i . On input j , Bob measures the state he receives with a POVM $\{E_\theta^j\}$ for some *nonnegative* values θ , and outputs the result. We say that such a protocol *computes a matrix M in expectation*, if the expected value of the output on respective inputs i and j , equals the matrix entry M_{ij} . Analogous to the equivalence between classical protocols and nonnegative factorizations of M established by Faenza et al. [26], such quantum protocols are essentially equivalent to PSD factorizations of S :

Theorem 17. *Let $M \in \mathbb{R}_+^{m \times n}$ be a matrix. Then the following holds:*

- (i) *A one-way quantum protocol with r -dimensional messages that computes M in expectation, gives a rank- r PSD factorization of M .*
- (ii) *A rank- r PSD factorization of M gives a one-way quantum protocol with $(r+1)$ -dimensional messages that computes M in expectation.*

Proof. The first part is straightforward. Given a quantum protocol as above, define $E^j := \sum_{\theta \in \Theta} \theta E_\theta^j$. Clearly, on inputs i and j the expected value of the output is $\text{Tr}[\rho_i E^j] = M_{ij}$.

For the second part, suppose we are given a PSD factorization of a matrix M , so we are given PSD matrices T_1, \dots, T_m and U^1, \dots, U^n satisfying $\text{Tr}[T_i U^j] = M_{ij}$ for all i, j . In order to turn this into a quantum protocol, define $\tau = \max_i \text{Tr}[T_i]$. Let ρ_i be the $(r+1)$ -dimensional quantum state obtained by adding a $(r+1)$ st row and column to T_i/τ , with $1 - \text{Tr}[T_i]/\tau$ as $(r+1)$ st diagonal entry, and 0s elsewhere. Note that ρ_i is indeed a PSD matrix of trace 1, so it is a well-defined quantum state. For input j , derive Bob's $(r+1)$ -dimensional POVM from the PSD matrix U^j as follows. Let λ be the largest eigenvalue of U^j , and define $E_{\tau\lambda}^j$ to be U^j/λ , extended with a $(d+1)$ st row and column of 0s. Let $E_0^j = I - E_{\tau\lambda}^j$. These two operators together form a well-defined POVM. The expected outcome (on inputs i, j) of the protocol induced by the states and POVMs that we just defined, is

$$\tau \lambda \text{Tr} \left[E_{\tau\lambda}^j \rho_i \right] = \text{Tr} [T_i U^j] = M_{ij},$$

so the protocol indeed computes M in expectation. □

We obtain the following corollary which summarizes the characterization of semidefinite EFs:

Corollary 18. *For a polytope P with slack matrix S , the following are equivalent:*

- (i) *P has a semidefinite extension $Q = \{(x, y) \in \mathbb{R}^{d+r(r+1)/2} \mid Ex + Fy = g, y \in \mathbb{S}_+^r\}$;*

⁷For simplicity we restrict to real rather than complex entries, which does not significantly affect the results.

- (ii) the slack matrix S has a rank- r PSD factorization;
- (iii) there exists a one-way quantum communication protocol with $(r + 1)$ -dimensional messages (i.e., using $\lceil \log(r + 1) \rceil$ qubits) that computes S in expectation (for the converse we consider r -dimensional messages).

5.2 A General Upper Bound on Quantum Communication

Now we provide a quantum protocol that efficiently computes a nonnegative matrix M in expectation, whenever there is a low rank matrix N whose entry-wise square is M .

Theorem 19. *Let M be a matrix with nonnegative real entries, N be a rank- r matrix of the same dimensions such that $M_{ij} = N_{ij}^2$. Then there exists a one-way quantum protocol using $(r + 1)$ -dimensional pure-state messages that computes M in expectation.*

Proof. By Corollary 18, it suffices to give a rank- r PSD factorization of M . To this end, let t_i, u_j be r -dimensional real vectors such that $N_{ij} = t_i \cdot u_j$; such vectors exist because N has rank r . Define $r \times r$ PSD matrices $T_i := t_i t_i^T$ and $U^j := u_j u_j^T$. Then

$$\text{Tr} [T_i U^j] = (t_i \cdot u_j)^2 = N_{ij}^2 = M_{ij},$$

hence we have a rank- r PSD factorization of M . □

Note that if M is a 0/1-matrix then we may take $N = M$, hence any low-rank 0/1-matrix can be computed in expectation by an efficient quantum protocol. If this M is the slack matrix for a polytope $P \subseteq \mathbb{R}^d$, then it is easy to see that its rank is at most $d + 1$: the slack $M_{ij} = b_i - A_i v_j$ of a constraint $A_i x \leq b_i$ w.r.t. a point $v_j \in P$ can be written as the inner product between the two $(d + 1)$ -dimensional vectors $(b_i, -A_i)$ and $(1, v_j)$. We thus obtain the following corollary (implicit in Theorem 4.2 of [34]) which also implies a compact (i.e., polynomial size) semidefinite EF for the stable set polytope of perfect graphs, reproving the previously known result by Lovász [52, 53]. We point out that the result still holds when $\dim(P) + 2$ is replaced by $\dim(P) + 1$, see [35]; this difference is due to normalization.

Corollary 20. *Let P be a polytope such that $S(P)$ is a 0/1 matrix. Then $\text{xc}_{SDP}(P) \leq \dim(P) + 2$.*

5.3 Quantum vs Classical Communication, and PSD vs Nonnegative Factorizations

We now give an example of an exponential separation between quantum and classical communication in expectation, based on the matrix M of Section 2. This result actually preceded and inspired the results in Section 3.

Theorem 21. *For each n , there exists a nonnegative matrix $M \in \mathbb{R}^{2^n \times 2^n}$ that can be computed in expectation by a quantum protocol using $\log n + O(1)$ qubits, while any classical randomized protocol needs $\Omega(n)$ bits to compute M in expectation.*

Proof. Consider the matrix $N \in \mathbb{R}^{2^n \times 2^n}$ whose rows and columns are indexed by n -bit strings a and b , respectively, and whose entries are defined as $N_{ab} = 1 - a^\top b$. Define $M \in \mathbb{R}_+^{2^n \times 2^n}$ by $M_{ab} = N_{ab}^2$. This M is the matrix from Section 2. Note that N has rank $r \leq n + 1$ because it can be written as the sum of $n + 1$ rank-1 matrices. Hence Theorem 19 immediately implies a quantum protocol with $(n + 2)$ -dimensional messages that computes M in expectation.

For the classical lower bound, note that a protocol that computes M in expectation has positive probability of giving a nonzero output on input a, b if and only if $M_{ab} > 0$. With a message m in this protocol we can associate a rectangle $R_m = A \times B$ where A consists of all inputs a for which Alice has positive probability of sending m , and B consists of all inputs b for which Bob, when he receives message m , has positive

probability of giving a nonzero output. Together these rectangles will cover exactly the nonzero entries of M . Accordingly, a c -bit protocol that computes M in expectation induces a rectangle cover for the support matrix of M of size 2^c . Theorem 1 lower bounds the size of such a cover by $2^{\Omega(n)}$, hence $c = \Omega(n)$. \square

Together with Theorem 17 and the equivalence of randomized communication complexity (in expectation) and nonnegative rank established in [26], we immediately obtain an exponential separation between nonnegative rank and PSD rank.

Corollary 22. *For each n , there exists $M \in \mathbb{R}_+^{2^n \times 2^n}$, with $\text{rank}_+(M) = 2^{\Omega(n)}$ and $\text{rank}_{PSD}(M) = O(n)$.*

In fact a simple rank- $(n+1)$ PSD factorization of M is the following: let $T_a := \begin{pmatrix} 1 \\ -a \end{pmatrix} \begin{pmatrix} 1 \\ -a \end{pmatrix}^\top$ and $U^b := \begin{pmatrix} 1 \\ b \end{pmatrix} \begin{pmatrix} 1 \\ b \end{pmatrix}^\top$, then $\text{Tr}[(T_a)^\top U^b] = (1 - a^\top b)^2 = M_{ab}$.

6 Concluding Remarks

In addition to proving the first unconditional super-polynomial lower bounds on the size of linear EFs for the cut polytope, stable set polytope and TSP polytope, we demonstrate that the rectangle covering bound can prove strong results in the context of EFs. In particular, it can be super-polynomial in the dimension and the logarithm of the number of vertices of the polytope, settling an open problem of [29].

The exponential separation between nonnegative rank and PSD rank that we prove here (Theorem 21) actually implies more than a super-polynomial lower bound on the extension complexity of the cut polytope. As noted in Theorem 5, the polytopes $\text{CUT}(n)$ and $\text{COR}(n-1)$ are affinely isomorphic. Let $Q(n)$ denote the polyhedron isomorphic (under the same affine map) to the polyhedron defined by (4) for $a \in \{0, 1\}^n$. Then (i) every polytope (or polyhedron) that contains $\text{CUT}(n)$ and is contained in $Q(n)$ has exponential extension complexity; (ii) there exists a low complexity spectrahedron that contains $\text{CUT}(n)$ and is contained in $Q(n)$. (A *spectrahedron* is any projection of an affine slice of the positive semidefinite cone.)

An important problem also left open in [72] is whether the perfect matching polytope has a polynomial-size linear EF. Yannakakis proved that every *symmetric* EF of this polytope has exponential size, a striking result given the fact that the perfect matching problem is solvable in polynomial time. He conjectured that asymmetry also does not help in the case of the perfect matching polytope. Because it is based on the rectangle covering bound, our argument does not yield a super-polynomial lower bound on the extension complexity of the perfect matching polytope. However, when we were revising the current paper, Rothvoß[62] posted on arXiv a proof of the fact that the extension complexity of the perfect matching polytope is $2^{\Omega(n)}$. This groundbreaking result is based on a general lower bound called the *hyperplane separation bound*, which was used implicitly, e.g., in [10].

As mentioned at the end of the introduction, the new connections developed have already inspired much follow-up research in particular about *approximate* EFs. Here are two concrete questions left open for future work: (i) find a *slack matrix* that has an exponential gap between nonnegative rank and PSD rank; (ii) prove that the cut polytope has no polynomial-size *semidefinite* EF (that would rule out SDP-based algorithms for optimizing over the cut polytope, in the same way that this paper ruled out LP-based algorithms).

Our final remark concerns the famous *log-rank conjecture* [54]. It states that the deterministic communication complexity of a (finite) Boolean matrix M is upper bounded by a polynomial in the logarithm of its rank $\text{rank}(M)$. On the one hand, this conjecture is equivalent to the following statement: $\log(\text{rank}_+(M)) \leq \text{poly}(\log(\text{rank}(M)))$ for all Boolean matrices M . On the other hand, we know that $\text{rank}_{PSD}(M) = O(\text{rank}(M))$ for all Boolean matrices M by Theorem 19. Using the interpretation of the nonnegative and PSD rank of M in terms of classical and quantum communication protocols computing M in expectation (see [26] and Theorem 17), we see that the log-rank conjecture is *equivalent* to the conjecture that classical protocols computing M in expectation are at most polynomially less efficient than quantum protocols. Accordingly, one way to prove the log-rank conjecture would be to give an efficient classical simulation of such quantum protocols for Boolean M (for *non-Boolean* M , we already exhibited an exponential separation in this paper).

Acknowledgments

We thank Monique Laurent for information about hypermetric inequalities, and the three anonymous STOC'12 referees as well as one JACM referee for suggesting improvements to the text. Sebastian Pokutta would like to thank Alexander Martin for the inspiring discussions and support. Ronald de Wolf thanks Giannicola Scarpa and Troy Lee for useful discussions.

Samuel Fiorini acknowledges support from the *Actions de Recherche Concertées* (ARC) fund of the French community of Belgium. Serge Massar acknowledges support from the European Commission under the projects QCS (Grant No. 255961) and QALGO (Grant No. 600700). Hans Raj Tiwary was postdoctoral researcher of the *Fonds National de la Recherche Scientifique* (F.R.S.–FNRS). Ronald de Wolf was partially supported by a Vidi grant from the Netherlands Organization for Scientific Research (NWO), and by the European Commission under the projects QCS (Grant No. 255961) and QALGO (Grant No. 600700).

References

- [1] S. Aaronson. Lower bounds for local search by quantum arguments. *SIAM Journal on Computing*, 35(4):804–824, 2006. Earlier version in STOC'04.
- [2] D. Aharonov and O. Regev. Lattice problems in $\text{NP} \cap \text{coNP}$. In *Proc. FOCS 2004*, pages 362–371, 2004.
- [3] S. Arora, B. Bollobás, and L. Lovász. Proving integrality gaps without knowing the linear program. In *Proc. FOCS 2002*, pages 313–322, 2002.
- [4] S. Arora, B. Bollobás, L. Lovász, and I. Tourlakis. Proving integrality gaps without knowing the linear program. *Theory Comput.*, 2:19–51, 2006.
- [5] D. Avis and H. R. Tiwary. On the extension complexity of combinatorial polytopes. In *Proc. ICALP(1) 2013*, pages 57–68, 2013.
- [6] E. Balas. Disjunctive programming and a hierarchy of relaxations for discrete optimization problems. *SIAM J. Algebraic Discrete Methods*, 6:466–486, 1985.
- [7] E. Balas, S. Ceria, and G. Cornuéjols. A lift-and-project algorithm for mixed 0-1 programs. *Math. Programming*, 58:295–324, 1993.
- [8] S. Benabbas and A. Magen. Extending SDP integrality gaps to Sherali-Adams with applications to quadratic programming and MaxCutGain. In *Proc. IPCO 2010*, pages 299–312, 2010.
- [9] G. Braun, S. Fiorini, and S. Pokutta. Average case polyhedral complexity of the maximum stable set problem. *arXiv:1311.4001*, 2013.
- [10] G. Braun, S. Fiorini, S. Pokutta, and D. Steurer. Approximation Limits of Linear Programs (Beyond Hierarchies). In *Proc. FOCS 2012*, pages 480–489, 2012.
- [11] G. Braun, R. Jain, T. Lee, and S. Pokutta. Information-theoretic approximations of the nonnegative rank. *ECCC:2013-158*, 2013.
- [12] G. Braun and S. Pokutta. Common information and unique disjointness. In *Proc. FOCS 2013*, 2013.
- [13] M. Braverman and A. Moitra. An information complexity approach to extended formulations. In *Proc. STOC 2013*, pages 161–170, 2013.
- [14] J. Briët, D. Dadush, and S. Pokutta. On the existence of 0/1 polytopes with high semidefinite extension complexity, 2013.

- [15] H. Buhrman, R. Cleve, S. Massar, and R. de Wolf. Nonlocality and communication complexity. *Rev. Modern Phys.*, 82:665, 2010.
- [16] J. Buresh-Oppenheim, N. Galesi, S. Hoory, A. Magen, and T. Pitassi. Rank bounds and integrality gaps for cutting planes procedures. *Theory Comput.*, 2:65–90, 2006.
- [17] S.O. Chan, J.R. Lee, P. Raghavendra, and D. Steurer. Approximate constraint satisfaction requires large LP relaxations. FOCS’13, 2013.
- [18] M. Charikar, K. Makarychev, and Y. Makarychev. Integrality gaps for Sherali-Adams relaxations. In *Proc. STOC 2009*, pages 283–292, 2009.
- [19] M. Conforti, G. Cornuéjols, and G. Zambelli. Extended formulations in combinatorial optimization. *4OR*, 8:1–48, 2010.
- [20] G.B. Dantzig. Maximization of a linear function of variables subject to linear inequalities. In *Activity Analysis of Production and Allocation*, Cowles Commission Monograph No. 13, pages 339–347. John Wiley & Sons Inc., New York, N. Y., 1951.
- [21] C. De Simone. The cut polytope and the Boolean quadric polytope. *Discrete Math.*, 79:71–75, 1990.
- [22] R. de Wolf. Quantum communication and complexity. *Theoret. Comput. Sci.*, 287:337–353, 2002.
- [23] R. de Wolf. Nondeterministic quantum query and communication complexities. *SIAM J. Comput.*, 32:681–699, 2003.
- [24] M.M. Deza and M. Laurent. *Geometry of cuts and metrics*, volume 15 of *Algorithms and Combinatorics*. Springer-Verlag, 1997.
- [25] A. Drucker and R. de Wolf. Quantum proofs for classical theorems. *Theory of Computing, Graduate Surveys*, 2, 2011.
- [26] Y. Faenza, S. Fiorini, R. Grappe, and H. R. Tiwary. Extended formulations, non-negative factorizations and randomized communication protocols. arXiv:1105.4127, 2011.
- [27] H. Fawzi and P.A. Parrilo. Exponential lower bounds on fixed-size psd rank and semidefinite extension complexity. arXiv:1311.2571, 2013.
- [28] W. Fernandez de la Vega and C. Mathieu. Linear programming relaxation of Maxcut. In *Proc. SODA 2007*, pages 53–61, 2007.
- [29] S. Fiorini, V. Kaibel, K. Pashkovich, and D. O. Theis. Combinatorial bounds on nonnegative rank and extended formulations. arXiv:1111.0444, 2011.
- [30] S. Fiorini, S. Massar, M. K. Patra, and H. R. Tiwary. Generalised probabilistic theories and conic extensions of polytopes, 2013. CoRR abs/1310.4125.
- [31] K. Georgiou, A. Magen, T. Pitassi, and I. Tzourakis. Integrality gaps of $2 - o(1)$ for vertex cover SDPs in the Lovász-Schrijver hierarchy. *SIAM J. Comput.*, 39:3553–3570, 2010.
- [32] K. Georgiou, A. Magen, and M. Tulsiani. Optimal Sherali-Adams gaps from pairwise independence. In *Proc. APPROX-RANDOM 2009*, pages 125–139, 2009.
- [33] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42:1115–1145, 1995.
- [34] J. Gouveia, P.A. Parrilo, and R. Thomas. Theta bodies for polynomial ideals. *SIAM J. Optim.*, 20:2097–2118, 2010.

- [35] J. Gouveia, P.A. Parrilo, and R. Thomas. Lifts of convex sets and cone factorizations. arXiv:1111.3164, 2011.
- [36] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999. Earlier version in FOCS’96.
- [37] H. Huang and B. Sudakov. A counterexample to the Alon-Saks-Seymour conjecture and related problems. *Combinatorica*, 32(2):205–219, 2012.
- [38] R. Jain, Y. Shi, Z. Wei, and S. Zhang. Efficient protocols for generating bipartite classical distributions and quantum states. *IEEE Transactions on Information Theory*, 59(8):5171–5178, 2013.
- [39] V. Kaibel. Extended formulations in combinatorial optimization. *Optima*, 85:2–7, 2011.
- [40] V. Kaibel, K. Pashkovich, and D.O. Theis. Symmetry matters for the sizes of extended formulations. In *Proc. IPCO 2010*, pages 135–148, 2010.
- [41] V. Kaibel and S. Weltge. A short proof that the extension complexity of the correlation polytope grows exponentially, 2013. arXiv:1307.3543.
- [42] N. Karmakar. A new polynomial time algorithm for linear programming. *Combinatorica*, 4:373–395, 1984.
- [43] I. Kerenidis and R. de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. *Journal of Computer and System Sciences*, 69(3):395–420, 2004. Earlier version in STOC’03.
- [44] L. G. Khachiyan. A polynomial algorithm in linear programming. *Dokl. Akad. Nauk SSSR*, 244(5):1093–1096, 1979.
- [45] S. Khot, G. Kindler, E. Mossel, and R. O’Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM J. Comput.*, 37(1):319–357, 2007.
- [46] H. Klauck, T. Lee, and S. Zhang. An explicit and exponential separation between randomized and quantum correlation complexities, 2011. Unpublished manuscript from Oct/Nov 2011. Personal communication between Troy Lee and Ronald de Wolf, December 2011 at Centre for Quantum Technologies, Singapore.
- [47] E. Kushilevitz and N. Nisan. *Communication complexity*. Cambridge University Press, 1997.
- [48] E. Kushilevitz and E. Weinreb. The communication complexity of set-disjointness with small sets and 0-1 intersection. In *Proc. FOCS 2009*, pages 63–72, 2009.
- [49] E. Kushilevitz and E. Weinreb. On the complexity of communication complexity. In *Proc. STOC 2009*, pages 465–474, 2009.
- [50] T. Lee and D. O. Theis. Support-based lower bounds for the positive semidefinite rank of a nonnegative matrix, 2012. arXiv:1203.3961.
- [51] C. Lemaréchal and J.B. Hiriart-Urruty. *Convex analysis and minimization algorithms I*. Springer, 1996.
- [52] L. Lovász. On the Shannon capacity of a graph. *IEEE Trans. Inform. Theory*, 25:1–7, 1979.
- [53] L. Lovász. Semidefinite programs and combinatorial optimization. In *Recent advances in algorithms and combinatorics*, volume 11 of *CMS Books Math./Ouvrages Math. SMC*, pages 137–194. Springer, 2003.
- [54] L. Lovász and M. Saks. Communication complexity and combinatorial lattice theory. *J. Comput. System Sci.*, 47:322–349, 1993.

- [55] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM J. Optim.*, 1:166–190, 1991.
- [56] N.D. Mermin. *Quantum computer science: an introduction*. Cambridge University Press, 2007.
- [57] M. A. Nielsen and I. L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2000.
- [58] K. Pashkovich. Symmetry in extended formulations of the permutahedron. arXiv:0912.3446, 2009.
- [59] S. Pokutta and M. Van Vyve. A note on the extension complexity of the knapsack polytope. *Operations Research Letters*, 41(4):347–350, 2013.
- [60] A. A. Razborov. On the distributional complexity of disjointness. *Theoret. Comput. Sci.*, 106(2):385–390, 1992.
- [61] T. Rothvoß. Some 0/1 polytopes need exponential size extended formulations. arXiv:1105.0036, 2011.
- [62] T. Rothvoß. The matching polytope has exponential extension complexity. arXiv:1311.2369, 2013.
- [63] G. Schoenebeck, L. Trevisan, and M. Tulsiani. Tight integrality gaps for Lovasz-Schrijver LP relaxations of vertex cover and max cut. In *Proc. STOC 2007*, pages 302–310, 2007.
- [64] A. Schrijver. *Combinatorial optimization. Polyhedra and efficiency*. Springer-Verlag, 2003.
- [65] C. E. Shannon. The synthesis of two-terminal switching circuits. *Bell System Tech. J.*, 25:59–98, 1949.
- [66] H. D. Sherali and W. P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM J. Discrete Math.*, 3:411–430, 1990.
- [67] M. Sipser. *Introduction to the Theory of Computation*. PWS, Boston, MA, 1996.
- [68] E. R. Swart. P = NP. Technical report, University of Guelph, 1986; revision 1987.
- [69] F. Vanderbeck and L. A. Wolsey. Reformulation and decomposition of integer programs. In M. Jünger et al., editor, *50 Years of Integer Programming 1958-2008*, pages 431–502. Springer, 2010.
- [70] L. A. Wolsey. Using extended formulations in practice. *Optima*, 85:7–9, 2011.
- [71] M. Yannakakis. Expressing combinatorial optimization problems by linear programs (extended abstract). In *Proc. STOC 1988*, pages 223–228, 1988.
- [72] M. Yannakakis. Expressing combinatorial optimization problems by linear programs. *J. Comput. System Sci.*, 43(3):441–466, 1991.
- [73] G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, 1995.

A Background on Polytopes

A (*convex*) *polytope* is a set $P \subseteq \mathbb{R}^d$ that is the convex hull $\text{conv}(V)$ of a finite set of points V . Equivalently, P is a polytope if and only if P is bounded and the intersection of a finite collection of closed halfspaces. This is equivalent to saying that P is bounded and the set of solutions of a finite system of linear inequalities and possibly equalities (each of which can be represented by a pair of inequalities).

Let $P \subseteq \mathbb{R}^d$ be a polytope. A closed halfspace H^+ that contains P is said to be *valid* for P . In this case the hyperplane H that bounds H^+ is also said to be *valid* for P . A *face* of P is either P itself or the intersection of P with a valid hyperplane. Every face of a polytope is again a polytope. A face is called

proper if it is not the polytope itself. A *vertex* is a minimal nonempty face. A *facet* is a maximal proper face. An inequality $c^\top x \leq \delta$ is said to be *valid* for P if it is satisfied by all points of P . The face it defines is $F := \{x \in P \mid c^\top x = \delta\}$. The inequality is called *facet-defining* if F is a facet. The *dimension* of a polytope P is the dimension of its affine hull $\text{aff}(P)$.

Every (finite or infinite) set V such that $P = \text{conv}(V)$ contains all the vertices of P . Conversely, letting $\text{vert}(P)$ denote the vertex set of P , we have $P = \text{conv}(\text{vert}(P))$. Suppose now that P is *full-dimensional*, i.e., $\dim(P) = d$. Then, every (finite) system $Ax \leq b$ such that $P = \{x \in \mathbb{R}^d \mid Ax \leq b\}$ contains all the facet-defining inequalities of P , up to scaling by positive numbers. Conversely, P is described by its facet-defining inequalities.

If P is not full-dimensional, these statements have to be adapted as follows. Every (finite) system describing P contains all the facet-defining inequalities of P , up to scaling by positive numbers and adding an inequality that is satisfied with equality by *all* points of P . Conversely, a linear description of P can be obtained by picking one inequality per facet and adding a system of equalities describing $\text{aff}(P)$.

A *0/1-polytope* in \mathbb{R}^d is simply the convex hull of a subset of $\{0, 1\}^d$.

A (convex) *polyhedron* is a set $P \subseteq \mathbb{R}^d$ that is the intersection of a finite collection of closed halfspaces. A polyhedron P is a polytope if and only if it is bounded.

For more background on polytopes and polyhedra, see the standard reference [73].