From Trees to Polynomials and Back Again: New Capacity Bounds with Applications to TSP

Leonid Gurvits
City College New York, NY, USA

Nathan Klein
Institute for Advanced Study, Princeton, NJ, USA

Jonathan Leake
University of Waterloo, Canada

Abstract

We give simply exponential lower bounds on the probabilities of a given strongly Rayleigh distribution, depending only on its expectation. This resolves a weak version of a problem left open by Karlin-Klein-Oveis Gharan in their recent breakthrough work on metric TSP, and this resolution leads to a minor improvement of their approximation factor for metric TSP. Our results also allow for a more streamlined analysis of the algorithm.

To achieve these new bounds, we build upon the work of Gurvits-Leake on the use of the productization technique for bounding the capacity of a real stable polynomial. This technique allows one to reduce certain inequalities for real stable polynomials to products of affine linear forms, which have an underlying matrix structure. In this paper, we push this technique further by characterizing the worst-case polynomials via bipartitioned forests. This rigid combinatorial structure yields a clean induction argument, which implies our stronger bounds.

In general, we believe the results of this paper will lead to further improvement and simplification of the analysis of various combinatorial and probabilistic bounds and algorithms.

2012 ACM Subject Classification Mathematics of computing → Approximation algorithms; Mathematics of computing → Probability and statistics; Theory of computation → Routing and network design problems

Keywords and phrases traveling salesman problem, strongly Rayleigh distributions, polynomial capacity, probability lower bounds, combinatorial lower bounds

Digital Object Identifier 10.4230/LIPIcs.ICALP.2024.79

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2311.09072`

Funding *Nathan Klein*: This author is supported by the National Science Foundation under Grant No. DMS-1926686.
Jonathan Leake: This author acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), [funding reference number RGPIN-2023-03726]. Cette recherche a été partiellement financée par le Conseil de recherches en sciences naturelles et en génie du Canada (CRSNG), [numéro de référence RGPIN-2023-03726].

Acknowledgements The third author would like to thank Vijay Bhattiprolu for many helpful and interesting discussions.

1 Introduction

The theory of real stable and log-concave polynomials has seen numerous applications in combinatorics and theoretical computer science (TCS). This includes bounds and approximation algorithms for various combinatorial quantities [18, 11, 7, 28, 5, 8, 2, 15], proofs of long-standing log-concavity and sampling conjectures related to matroids [1, 3, 4, 14], proofs of the Kadison-Singer conjecture and generalizations [26, 6, 13], an improved approximation factor for the traveling salesperson problem (TSP) [27, 24, 25], and many more. The power of these polynomial classes comes from two features: (1) their *robustness*, shown in the fact that many natural operations preserve these log-concavity properties, and (2) their *convex analytic properties*, which can be used to prove bounds and other analytic statements. The typical way these polynomials are utilized is by encoding combinatorial objects as real stable and log-concave polynomials, which essentially allows these operations and convexity properties to automatically transfer to the combinatorial objects. This idea, while simple, has led to important breakthroughs in combinatorics, TCS, and beyond.

For example, [18] utilized real stable polynomials to give a new proof of the Van der Waerden conjecture on the permanent of a doubly stochastic matrix (originally due to [16, 17]). This proof led to a vast generalization of Van der Waerden bound, including an improved Schrijver's bound for regular bipartite graphs [18], an analogous bound for mixed discriminants [19], and an analogous bound for mixed volumes that led to the development of strongly log-concave polynomials [21, 20]. One reason the original bound was historically so difficult to prove is a lack of a usable inductive structure coming from the matrices themselves. One of the key insights of the new proof was to use the simple inductive structure of real stable polynomials given by partial derivatives. By encoding the matrices as polynomials, the correct induction becomes straightforward, and the bound follows from a simple calculus argument.

More recently, the approximation factor improvement for the metric traveling salesperson problem (TSP) crucially utilized real stable polynomials [24, 25]. The idea is to encode certain discrete probability distributions related to spanning trees as real stable polynomials. The coefficients of these polynomials give probabilities of certain graph-theoretic events (e.g., the number of edges in a given spanning tree incident on a particular vertex), and analytic properties of real stable polynomials allow one to lower bound these probabilities. This in turn implies bounds on the expected cost of the output of a randomized algorithm for metric TSP.

In this paper, we improve upon the polynomial capacity bounds of [22], and our applications touch on the two problems discussed above. Specifically, we give:

1. robust coefficient lower bounds for all (not necessarily homogeneous) real stable polynomials,
2. simply exponential lower bounds on probabilities of strongly Rayleigh distributions (solving a weak version of an open problem of [24]), and
3. a further improvement to the approximation factor for metric TSP (predicted by [22]).

Interestingly, our approach goes in the opposite direction to that discussed above. Our technical results answer questions regarding real stable polynomials, but to prove these results we use various graph and matrix structures inherent to the polynomials. In [22], this was seen in the "productization" technique: bounds on real stable polynomials were achieved by showing that the worst-case bounds come from polynomials associated to certain matrices. In this paper, we push this idea further by showing that these worst-case matrices are bipartite adjacency matrices of forests. This very rigid structure enables a clean induction argument, which implies stronger polynomial capacity bounds. These new bounds lead to the applications discussed above, with the strongest bounds implying the metric TSP improvement.

2 Main Results

We first state here our main technical results; see Section 3 for any undefined notation.

Our first main result is a non-homogeneous version of Theorem 2.1 of [22] which implies robust coefficient lower bounds for all real stable polynomials as a direct corollary. Crucially, these bounds do not depend on the total degree of the polynomial. This was one of the main barriers to applying the results of [22] to metric TSP.

▶ **Theorem 1** (Main non-homogeneous capacity bound). *Let $p \in \mathbb{R}_{\geq 0}[x_1, \ldots, x_n]$ be a real stable polynomial in n variables, and fix any $\boldsymbol{\kappa} \in \mathbb{Z}^n$ with non-negative entries. If $p(\mathbf{1}) = 1$ and $\|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1 < 1$, then*

$$\inf_{x_1,\ldots,x_n > 0} \frac{p(\boldsymbol{x})}{x_1^{\kappa_1} \cdots x_n^{\kappa_n}} \geq \left(1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1\right)^n .$$

This bound is tight for any fixed $\boldsymbol{\kappa}$ with strictly positive entries.

▶ **Corollary 2** (Main non-homogeneous coefficient bound). *Let $p \in \mathbb{R}_{\geq 0}[x_1, \ldots, x_n]$ be a real stable polynomial in n variables, and fix any $\boldsymbol{\kappa} \in \mathbb{Z}^n$ with non-negative entries. If $p(\mathbf{1}) = 1$ and $\|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1 < 1$, then*

$$p_{\boldsymbol{\kappa}} \geq \left(\prod_{i=1}^{n} \frac{\kappa_i^{\kappa_i} e^{-\kappa_i}}{\kappa_i!}\right) (1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)^n,$$

where $p_{\boldsymbol{\kappa}}$ is the coefficient of $\boldsymbol{x}^{\boldsymbol{\kappa}}$ in p. The dependence on $(1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)$ is tight for any fixed $\boldsymbol{\kappa}$ with strictly positive entries.

The above results[1] are robust (i.e., resilient to ℓ^1 perturbations) versions of the results utilized to bound various combinatorial and probabilistic quantities, as discussed above. That said, they are still not quite strong enough to imply an improvement to the metric TSP approximation factor. To obtain this improvement, we resolve a weak version of an open problem from [24], which we discuss below. Stronger versions of Theorem 1 and Corollary 2 which imply this result can be found in Section 5.

2.1 Application: Minimum Permanent

Before discussing the application to TSP, we first describe a different application of our bounds as a sort of prelude. It is at this point well-known that the permanent of any $n \times n$ doubly stochastic matrix is at least $\frac{n!}{n^n}$, and that $\frac{1}{n}\mathbf{1} \cdot \mathbf{1}^\top$ is the unique minimizer of the permanent over all doubly stochastic matrices. On the other hand, a similar tight lower bound with explicit minimizer is not known for sets of matrices with different row and column sums. The following then slightly extends what is known in the doubly stochastic case. See Section 6 for further details.

Given $\boldsymbol{c} \in \mathbb{R}^n_{\geq 0}$, let $\mathrm{Mat}_n(\boldsymbol{c})$ denote the set of $n \times n$ matrices with non-negative entries, row sums equal to $\mathbf{1}$, and columns sums equal to $\boldsymbol{c}$.

▶ **Theorem 3.** *For all $n \geq 1$, there exists $\epsilon_n > 0$ such that if $\|\boldsymbol{c} - \mathbf{1}\|_1 < \epsilon_n$ then $\frac{1}{n}\mathbf{1} \cdot \boldsymbol{c}^\top$ is the unique minimizer of the permanent over $\mathrm{Mat}_n(\boldsymbol{c})$. Specifically, this holds if $\frac{c_1 c_2 \cdots c_n}{L(\boldsymbol{c})} < \frac{(n-2)^{n-2} n^{n-1}}{(n-1)^{2n-3}}$, where $L(\boldsymbol{c})$ is any lower bound on the capacity.*

[1] Note that these bounds already appear in the arXiv version of [22], but not in the STOC version.

The proof utilizes Theorem 1 or results of [22] to guarantee all minimizers of the permanent lie in the relative interior of $\mathrm{Mat}_n(\boldsymbol{c})$. The symmetry and multilinearity of the permanent then imply the minimizer must be the rank-one matrix of $\mathrm{Mat}_n(\boldsymbol{c})$ described above. See Section Section 6 for the explicit value of ϵ_n and the details of the proof. It should be noted that in proving uniqueness, we were able to avoid usage of the conditions for equality in the Alexandrov-Fenchel inequalities.

On the other hand, when $\boldsymbol{c}$ is far from $\mathbf{1}$, the above result can be far from correct. Recall from [22] that $\mathrm{per}(M) > 0$ for all $M \in \mathrm{Mat}_n(\boldsymbol{c})$ if and only if $\|\boldsymbol{c} - \mathbf{1}\|_1 < 2$.

▶ **Proposition 4.** *For all n large enough, there exists $\boldsymbol{c}$ such that $\|\boldsymbol{c} - \mathbf{1}\|_1 < 2$ and a sparse matrix $M \in \mathrm{Mat}_n(\boldsymbol{c})$ (with linearly many non-zero entries) which has smaller permanent than that of $\frac{1}{n}\mathbf{1} \cdot \boldsymbol{c}^\top$.*

These two results suggest the complexity of the minimizer of the permanent, given the column sums $\boldsymbol{c}$. The coefficient bound given above in Corollary 2 is then a lower bound which generalizes that of the permanent to coefficients of real stable polynomials. (Consider the coefficient $p_{\mathbf{1}}$ of $p(\boldsymbol{x}) = \prod_{i=1}^n \sum_{j=1}^n m_{ij}x_j$.) Thus Corollary 2 can be seen as a sort of "smoothing" of the complexities that can occur for minima of the permanent and its generalizations.

Further, Theorem 3 can generalized to a permanent-like function on rectangular matrices using Theorem 1, and even beyond that to the mixed discriminant. (For the mixed discriminant, the row sum condition becomes a trace condition, and the column sum condition becomes an eigenvalue condition.) On the other hand, we can generalize the statement of Theorem 3 to coefficients of real stable polynomials in general, but we do not yet know how to prove it.

2.2 Application: Metric TSP

We first recall an important probabilistic bound from [24] used in the analysis of their metric TSP approximation algorithm (which is a slight modification of the max entropy algorithm from [27] first studied by [9]). In what follows, we let $A_S := \sum_{i \in S} A_i$ and $\kappa_S := \sum_{i \in S} \kappa_i$. See Section 3 for any undefined notation.

▶ **Theorem 5** (Prop. 5.1 of [24]). *Let μ be a strongly Rayleigh distribution on $[m]$, let $A_1, \ldots, A_n$ be random variables counting the number of elements contained in disjoint subsets of $[m]$, and fix $\boldsymbol{\kappa} \in \mathbb{Z}^n$ with non-negative entries. Suppose for all $S \subseteq [n]$ we have*

$$\mathbb{P}_\mu \left[A_S \geq \kappa_S\right] \geq \epsilon \qquad \text{and} \qquad \mathbb{P}_\mu \left[A_S \leq \kappa_S\right] \geq \epsilon.$$

Then we have

$$\mathbb{P}_\mu \left[A_1 = \kappa_1, \ldots, A_n = \kappa_n\right] \geq f(\epsilon) \cdot \mathbb{P}_\mu \left[A_{[n]} = \kappa_{[n]}\right],$$

where $f(\epsilon) \geq \epsilon^{2^n} \prod_{i=2}^n \frac{1}{\max\{\kappa_i, \kappa_{[i-1]}\}+1}$.

In [24], the authors note two things about this bound. First, they note that to apply the bound it is sufficient to have

$$\left|\mathbb{E}_\mu \left[A_S\right] - \kappa_S\right| < 1 \qquad \forall S \subseteq [n], \tag{1}$$

since this implies a lower bound on $\mathbb{P}_\mu \left[A_S = \kappa_S\right]$ for all $S \subseteq [n]$ for strongly Rayleigh distributions. Second, they note that the bound on $f(\epsilon)$ is doubly exponential in n, but they expect the true dependency to only be simply exponential. They leave it as an open problem to determine a tight lower bound on $f(\epsilon)$.

In this paper, we further improve the metric TSP approximation factor by resolving a weak version of this open problem: we give a simply exponential lower bound which depends tightly on ϵ, under the stronger condition of (1). Concretely, we prove the following.

▶ **Theorem 6** (Improved probability lower bound)**.** *Let μ be a strongly Rayleigh distribution on $[m]$, let $A_1, \ldots, A_n$ be random variables counting the number of elements contained in disjoint subsets of $[m]$, and fix $\boldsymbol{\kappa} \in \mathbb{Z}^n$ with non-negative entries. Suppose for all $S \subseteq [n]$ we have*

$$\left|\mathbb{E}_\mu\left[A_S\right] - \kappa_S\right| \leq 1 - \epsilon.$$

Then we have

$$\mathbb{P}_\mu\left[A_1 = \kappa_1, \ldots, A_n = \kappa_n\right] \geq \epsilon^n \prod_{\kappa_i > 0} \frac{1}{e\sqrt{\kappa_i}}.$$

The dependence on ϵ is tight for any fixed $\boldsymbol{\kappa}$ with strictly positive entries.

In fact, we prove stronger versions of this result which more directly depend on the specific values of $(\mathbb{E}_\mu\left[A_S\right] - \kappa_S)$ for all $S \subseteq [n]$; see Theorem 9 and Corollary 10. These results are analogous to our main coefficient bound Corollary 2 because the coefficients and the gradient of probability generating polynomials can be interpreted as the probabilities and the expectation of the associated distribution. That said, our stronger probabilistic results require a more delicate analysis of the expectations (gradient) beyond what is required for Corollary 2. In particular, note that the conditions on the expectations in Theorem 6 are more general than a bound on the ℓ^1 norm of $(\mathbb{E}_\mu\left[A_i\right] - \kappa_i)_{i=1}^n$. See Section 4 for further details.

Using Theorem 6, we improve the metric TSP approximation factor for the algorithm given in [24].

▶ **Theorem 7.** *There exists a randomized algorithm for metric TSP with approximation factor $\frac{3}{2} - \epsilon$ for some $\epsilon > 10^{-34}$.*

This is about a 100 times improvement over the result of [23]. Thus our improvement in terms of the approximation factor itself may be smaller than anticipated, given that we were able to improve the probability bound in Theorem 6 from doubly exponential to simply exponential. The reason for this is that while Theorem 5 was useful in [24] to quickly determine which events occurred with constant probability (and indeed provided a single unifying explanation for why one should expect many of their probabilistic bounds to hold), it gave such small guarantees that [24] resorted to ad hoc arguments instead to give their final probabilistic bounds.

We show that Theorem 6 alone can be used to give bounds that are comparable to the ad hod methods of [24] (and, in several important cases, much better) whenever the bounds came purely from information on the expectations as in (1). Thus, we believe our main contribution to work on metric TSP is a version of Theorem 5 that is "reasonable" to use, allowing one to show a similar approximation factor but with a more streamlined proof.

Unfortunately, not all of the bounds in [24] follow from expectation information, and two of them become bottlenecks for improving the approximation factor after applying Theorem 6 to the other statements. Thus, to demonstrate Theorem 7 we need to sharpen these bounds using other techniques. For one of these lemmas we show that the existing proof in [24] was far from tight, and in the other we refine their proof. In particular, using Theorem 6, we show we can reduce this second lemma to a special case that is possible to analyze more carefully. See Section 4 for further details.

3 Technical Overview

In this section we discuss the proof strategy of our main capacity and coefficient bounds, Theorem 1 and Corollary 2, and their stronger forms.

Notation

Given a vector $\boldsymbol{z} \in \mathbb{R}^E$ and a subset S of E, let $\boldsymbol{z}^S := \prod_{e \in S} z_e$. Let $\mu : \{0,1\}^E \to \mathbb{R}$ be a probability distribution over subsets of E. The generating polynomial $g_\mu \in \mathbb{R}_{\geq 0}[\{z_e\}_{e \in E}]$ of μ is defined as

$$g_\mu(z) := \sum_{S \subseteq E} \mu(S) \cdot \boldsymbol{z}^S.$$

The distribution μ is **strongly Rayleigh** if g_μ is real stable, where a polynomial $p \in \mathbb{R}[z_1, \dots, z_n]$ is **real stable** if $p(\boldsymbol{z}) \neq 0$ whenever $\Im(z_i) > 0$ for all $i \in [n]$ (i.e., when all inputs are in the complex upper half-plane). See [12] for much more on strongly Rayleigh measures. Further, given a polynomial $p \in \mathbb{R}_{\geq 0}[x_1, \dots, x_n]$ and $\boldsymbol{\kappa} \in \mathbb{Z}^n_{\geq 0}$, the **capacity** of p is defined as

$$\mathrm{Cap}_{\boldsymbol{\kappa}}(p) := \inf_{\boldsymbol{x} > 0} \frac{p(\boldsymbol{x})}{\boldsymbol{x}^{\boldsymbol{\kappa}}}.$$

Finally, we let $p_{\boldsymbol{\kappa}}$ denote the coefficient of $\boldsymbol{x}^{\boldsymbol{\kappa}}$ in p.

Conceptual strategy

We first give an overarching view of the strategy used to prove our main results, as well as the key similarities and differences compared to that of [22]. The general idea for proving our bounds is to find a simple and sparse underlying structure for the worst-case inputs. The space of all real stable polynomials can be complicated, but we show that the worst-case polynomials for our bounds are far simpler: they are "sparse" products of affine linear forms. More concretely, we reduce the space of input polynomials (and the corresponding combinatorial structures) as follows:

real stable polynomials	$\Longrightarrow$	products of linears	$\Longrightarrow$	sparse products of linears
matroids	$\Longrightarrow$	matrices	$\Longrightarrow$	forests

The first reduction step uses the idea of productization which was the key idea from [22]. This allows for one to utilize the matrix structure inherent to products of affine linear forms.

The second reduction step is then new to this paper. We first show that we may restrict to the extreme points of the set of matrices corresponding to products of affine linear forms, and then we show that these extreme matrices are supported on the edges of forests. This implies a significant decrease in density of the matrices: general graphs can have quadratically many edges, whereas forests can only have linearly many. This allows for an intricate but clean induction on the leaf vertices of these forests, which yields the strongest bounds of this paper. Additionally, it is this step that allows for bounds which do not depend on the total degree of the polynomial, and this was a crucial barrier to applying the bounds of [22] to metric TSP.

3.1 Conceptual Strategy, in More Detail

We now go through the steps of the conceptual strategy described above in more detail. Let us first restate our main capacity and coefficient bounds.

▶ **Theorem 8** (= Theorem 1 and Corollary 2)**.** *Let $p \in \mathbb{R}_{\geq 0}[x_1, \ldots, x_n]$ be a real stable polynomial in n variables, and fix any $\boldsymbol{\kappa} \in \mathbb{Z}^n$ with non-negative entries. If $p(\mathbf{1}) = 1$ and $\|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1 < 1$, then*

$$\mathrm{Cap}_{\boldsymbol{\kappa}}(p) \geq (1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)^n$$

and

$$p_{\boldsymbol{\kappa}} \geq \left(\prod_{i=1}^{n} \frac{\kappa_i^{\kappa_i} e^{-\kappa_i}}{\kappa_i!}\right) (1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)^n ,$$

*where $\mathrm{Cap}_{\boldsymbol{\kappa}}(p) := \inf_{\boldsymbol{x}>0} \frac{p(\boldsymbol{x})}{\boldsymbol{x}^{\boldsymbol{\kappa}}}$ is the **capacity** of p and $p_{\boldsymbol{\kappa}}$ is the coefficient of $\boldsymbol{x}^{\boldsymbol{\kappa}}$ in p. The dependence on $(1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)$ in these bounds is tight for any fixed $\boldsymbol{\kappa}$ with strictly positive entries.*

Analogous bounds required for the metric TSP application then follow from interpreting desired quantities as the coefficients and gradient of certain real stable polynomials. Specifically, the strongly Rayleigh probabilities we wish to lower bound are the coefficients of the corresponding real stable generating polynomial, and the expectations of the associated random variables are given by the gradient of that polynomial. We leave further details to Section 4.

We now discuss the proof of Theorem 8. First note that the coefficient bound follows from the capacity bound. This immediately follows from Corollary 3.6 of [20], which implies

$$p_{\boldsymbol{\kappa}} \geq \left(\prod_{i=1}^{n} \frac{\kappa_i^{\kappa_i} e^{-\kappa_i}}{\kappa_i!}\right) \mathrm{Cap}_{\boldsymbol{\kappa}}(p). \tag{2}$$

Thus what remains to be proven is the capacity bound

$$\mathrm{Cap}_{\boldsymbol{\kappa}}(p) \geq (1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)^n ,$$

which is precisely the bound of Theorem 1, as well as its tightness, which follows from considering a particular example of p (see Lemma 28).

The remainder of the proof then has four main steps. We also note here that these proof steps actually imply stronger bounds than Theorem 1, see Corollary 27 for the formal statement. These stronger bounds are required for the metric TSP application.

Step 1: Reduce to products of affine linear forms via productization

We first generalize the productization technique of [22] to non-homogeneous real stable polynomials. The upshot of this technique is that it implies it is sufficient to prove Theorem 1 for products of affine linear forms with non-negative coefficients (see Corollary 18). Such polynomials correspond to $d \times (n+1)$ $\mathbb{R}_{\geq 0}$-valued matrices with row sums $\mathbf{1}$ and column sums $\boldsymbol{\alpha}$ equal the entries of the gradient of the polynomial, via

$$\phi : A \mapsto \prod_{i=1}^{d} \left(a_{i,n+1} + \sum_{j=1}^{n} a_{i,j} x_j \right) .$$

This gives far more structure to work with, beyond that of real stable polynomials in general. This part is a straightforward generalization of the analogous result of [22].

Step 2: Reduce to extreme points

The set of $\mathbb{R}_{\geq 0}$-valued matrices with row sums $\mathbf{1}$ and column sums $\boldsymbol{\alpha}$ forms a convex polytope $P_{\boldsymbol{\alpha}}^d$, and thus the polynomials we now must consider correspond to the points of this polytope via the map ϕ defined above. Inspired by a result of Barvinok (see Lemma 19), we next show that the function

$$A \mapsto \mathrm{Cap}_{\kappa}(\phi(A))$$

is log-concave on the above described polytope. Since we want to minimize the capacity, this implies we may further restrict to the polynomials associated to the extreme points of the polytope.

Step 3: Extreme points correspond to bipartitioned forests

Any $\mathbb{R}_{\geq 0}$-valued matrix A can be interpreted as the weighted bipartite adjacency matrix of a bipartite graph, where the left vertices correspond to the rows of A and the right vertices correspond to the columns of A. A matrix $A \in P_{\boldsymbol{\alpha}}^d$ being extreme implies the associated bipartite graph has no cycles. This implies the associated bipartite graph is a forest (see Lemma 20). The sparsity properties of such matrices implies a simple structure for the associated polynomials, which is particularly amenable to an intricate but clean induction.[2]

Step 4: Induction on leaf vertices of the bipartitioned forests

Leaf vertices in the forest corresponding to a given matrix $A \in P_{\boldsymbol{\alpha}}^d$ indicate rows or columns of the matrix A which have exactly one non-zero entry. If a row of A has exactly one non-zero entry, the induction proceeds in a straightforward fashion, by simply removing the corresponding row of A and recalculating the column sums (see the $d \geq n+1$ case of the proof of Theorem 25).

If a column of A (say column i) has exactly one non-zero entry, then the induction is more complicated. We prove lemmas showing how much the capacity can change after applying the partial derivative ∂_{x_i} (when $\kappa_i \geq \alpha_i$, see Lemma 24) or setting x_i to 0 (when $\kappa_i < \alpha_i$, see Lemma 23). Since column i has only one entry, applying ∂_{x_i} corresponds to removing column i and the row of A which contains the non-zero entry, and setting x_i to 0 corresponds to removing column i. After renormalizing the row sums and recalculating the column sums, the proof again proceeds by induction. (See the proof of Theorem 25 to see the above arguments presented formally.) Example 29 and Example 30 show that the distinction between the $\kappa_i \geq \alpha_i$ and $\kappa_i < \alpha_i$ cases is not an artifact of the proof.

Some comments on tightness of the bounds

The main coefficient bound of Corollary 2 is proven via two different bounds, as discussed at the beginning of this section. That is, one first bounds the coefficient in terms of the capacity (2) via Corollary 3.6 of [20], and then one bounds the capacity via the steps outlined above. Thus while the capacity bound (Theorem 1) is tight for $\boldsymbol{\kappa} > \mathbf{0}$, the coefficient bound

[2] Note that it is already mentioned in [22] that the supports of the extreme points correspond to forests, but the application of this observation in [22] is somewhat "naive" and kind of brute-force: it was mainly used to describe an (inefficient) algorithm to compute the capacity lower bound for products of linear forms, which is not related to main lower bound in this paper. Additionally, its use in [22] is quite conceptually far from how it could actually be used to improve the TSP approximation factor.

(Corollary 2) may not be. We note that the coefficient bound is likely close to tight in the case that $(1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)$ is close to 1, but it seems this tightness deteriorates as $(1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)$ gets close to 0. That said, the dependence on $(1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)$ in Theorem 1 and Corollary 2 is tight for $\boldsymbol{\kappa} > \mathbf{0}$ by Lemma 28, though there is still potential for improvement in the more refined capacity bounds; see Section 5.5.

That said, tight coefficient lower bounds in the univariate case can be achieved by directly applying a bound of Hoeffding, and these lower bounds resemble the coefficient lower bound of Corollary 2. Thus one can view Corollary 2 as a step towards a multivariate generalization of Hoeffding's theorem. It is an interesting question whether or not the techniques used here can be extended to a full multivariate generalization of Hoeffding's theorem.

3.2 Example: The Univariate Case

In this section, we demonstrate the proof of Theorem 1 in the univariate case. This will serve as a sort of proof of concept for the more general proof.

Here we consider real stable non-homogeneous polynomials $p \in \mathbb{R}_{\geq 0}[x_1]$ such that $p(1) = 1$ and $\nabla p(1) = \alpha_1$, and we define $\epsilon := 1 - |\alpha_1 - \kappa_1| > 0$. By Step 1 above, we may assume that p is of the form

$$p(x_1) = \prod_{i=1}^{d} (a_{i,1}x_1 + a_{1,2}),$$

where A is an $\mathbb{R}_{\geq 0}$-valued $d \times 2$ matrix with row sums 1 and column sums $(\alpha_1, d - \alpha_1)$. By Steps 2-3, we may further assume A is the weighted bipartite adjacency matrix of a forest. If every row of A contains exactly one non-zero entry then $p(x_1) = x_1^{\kappa_1}$, and the result is trivial in this case. Otherwise, $d - 1$ rows of M have exactly one non-zero entry (see Lemma 21). Thus for some $k \leq d - 1$ we have

$$p(x_1) = x_1^k(ax_1 + b),$$

where $a + b = 1$ and $a + k = \alpha_1$. Since $\kappa_1 \in \mathbb{Z}_{\geq 0}$ and

$$\epsilon = 1 - |\alpha_1 - \kappa_1| = 1 - |a + k - \kappa_1|,$$

we have that k is equal to either κ_1 or $\kappa_1 - 1$. If $k = \kappa_1$, then

$$\mathrm{Cap}_{\kappa_1}(p) = \inf_{x_1 > 0} \frac{x_1^k(ax_1 + b)}{x_1^k} = b \qquad \text{and} \qquad \epsilon = 1 - |a + k - \kappa_1| = 1 - a = b.$$

If $k = \kappa_1 - 1$, then

$$\mathrm{Cap}_{\kappa_1}(p) = \inf_{x_1 > 0} \frac{x_1^k(ax_1 + b)}{x_1^{k+1}} = a \qquad \text{and} \qquad \epsilon = 1 - |a + k - \kappa_1| = 1 - (1 - a) = a.$$

Therefore in both cases we have

$$\mathrm{Cap}_{\kappa_1}(p) = \epsilon = (1 - |\alpha_1 - \kappa_1|),$$

which proves Theorem 1 in the univariate case and demonstrates the tight dependence on $(1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)$ in this case.

4 Proof of Application: Metric TSP

The following is the strongest probability lower bound, which we will use for the metric TSP application. In what follows, we let $A_S := \sum_{i \in S} A_i$ and $\kappa_S := \sum_{i \in S} \kappa_i$. Note that if

$$\left|\mathbb{E}_\mu\left[A_S\right] - \kappa_S\right| \leq 1 - \epsilon,$$

then Theorem 6 immediately follows from Theorem 9, Lemma 28, and a standard computation.

▶ **Theorem 9** (Strongest form of the probability bound)**.** *Let μ be a strongly Rayleigh distribution on $[m]$, let $A_1, \ldots, A_n$ be random variables counting the number of elements contained in some associated disjoint subsets of $[m]$, and fix $\boldsymbol{\kappa} \in \mathbb{Z}_{\geq 0}^n$. Suppose for all $S \subseteq [n]$ we have $|\mathbb{E}_\mu[A_S] - \kappa_S| < 1$. Define $\boldsymbol{\epsilon}, \boldsymbol{\delta} \in \mathbb{R}_{>0}^n$ via*

$$\delta_k := 1 + \min_{S \in \binom{[n]}{k}} \left(\mathbb{E}_\mu\left[A_S\right] - \kappa_S\right) \qquad \textit{and} \qquad \epsilon_k := 1 - \max_{S \in \binom{[n]}{k}} \left(\mathbb{E}_\mu\left[A_S\right] - \kappa_S\right).$$

Then we have

$$\mathbb{P}_\mu\left[A_1 = \kappa_1, \ldots, A_n = \kappa_n\right] \geq \prod_{i=1}^n \frac{\kappa_i^{\kappa_i} e^{-\kappa_i}}{\kappa_i!} \cdot \min_{0 \leq \ell \leq n} \prod_{k=1}^{\ell} \epsilon_k \prod_{k=1}^{n-\ell} \delta_k.$$

Proof. Let q be the probability generating polynomial of μ, and let p be the polynomial obtained by setting the variables of q associated to A_i to x_i for all $i \in [n]$, and setting all other variables to equal 1. Since μ is strongly Rayleigh, p is real stable. Further, $p(\mathbf{1}) = 1$, $\nabla p(\mathbf{1}) = (\mathbb{E}_\mu[A_i])_{i=1}^n$, and $\mathbb{P}_\mu[A_1 = \kappa_1, \ldots, A_n = \kappa_n]$ is the $\boldsymbol{x}^{\boldsymbol{\kappa}}$ coefficient of p. Thus Gurvits' capacity inequality (2) and Corollary 27 imply the desired result. ◀

We also give a slightly weaker bound which is a bit easier to use in practice. Note that Theorem 6 also follows from Corollary 10.

▶ **Corollary 10.** *Let μ be a strongly Rayleigh distribution on $[m]$, let $A_1, \ldots, A_n$ be random variables counting the number of elements contained in some associated disjoint subsets of $[m]$, and fix $\boldsymbol{\kappa} \in \mathbb{Z}_{\geq 0}^n$. Suppose for all $S \subseteq [n]$ we have $|\mathbb{E}_\mu[A_S] - \kappa_S| < 1$. Define $\boldsymbol{\epsilon} \in \mathbb{R}_{>0}^n$ via*

$$\epsilon_k := 1 - \max_{\substack{S \subseteq [n] \\ |S| \leq k}} \left|\mathbb{E}_\mu\left[A_S\right] - \kappa_S\right|$$

Then we have

$$\mathbb{P}_\mu\left[A_1 = \kappa_1, \ldots, A_n = \kappa_n\right] \geq \prod_{i=1}^n \frac{\kappa_i^{\kappa_i} e^{-\kappa_i}}{\kappa_i!} \cdot \prod_{k=1}^n \epsilon_k.$$

Proof. Follows from Theorem 9; see the proof of Corollary 26 for more details. ◀

The remainder of this section is devoted to demonstrating how one can use the above results to improve the approximation factor for metric TSP.

4.1 A Simple Application

We recall some definitions from [24]. There, we have a graph $G = (V, E)$ and a strongly Rayleigh (SR) distribution $\mu : 2^E \to \mathbb{R}_{\geq 0}$ supported on spanning trees of G. We let $x_e = \mathbb{P}_{T\sim\mu}[e \in T]$ and for a set of edges F let $x(F) = \sum_{e\in F} x_e$. Furthermore we let $\delta(S) = \{e = \{u, v\} \mid |e \cap S| = 1\}$. The guarantee on x is that $x \in P_{\text{Sub}}$,[3] where

$$P_{\text{Sub}} := \begin{cases} x(\delta(S)) \geq 2 & \forall\, S \subsetneq V \\ x(\delta(v)) = 2 & \forall v \in V \\ x_{\{u,v\}} \geq 0 & \forall\, u, v \in V \end{cases} \tag{3}$$

In the algorithm they analyze, one first samples a spanning tree T from μ and then add the minimum cost matching on the odd vertices of T. Their work involves analyzing the expected cost of this matching over the randomness of the sampled tree T. Unsurprisingly, the *parity* of vertices is therefore very important, as it determines which vertices are involved in the matching.

In this section, to give some sense of the utility of Theorem 9, we show an application in a simplified setting. Namely, we show that for any two vertices u, v, except in the special case that $x_{\{u,v\}} \approx \frac{1}{2}$, we have $\mathbb{P}_{T\sim\mu}[|\delta(u) \cap T| = |\delta(v) \cap T| = 2] \geq \Omega(1)$, i.e. for any two vertices that do not share an edge of value $\frac{1}{2}$ there is a constant probability that they have even parity simultaneously. This is helpful because (under some conditions on the point $x \in P_{\text{Sub}}$) the event that u, v are even simultaneously indicates one can strictly decrease the cost of the matching proportional to the cost of the edge e.

We now prove $\mathbb{P}_{T\sim\mu}[|\delta(u) \cap T| = |\delta(v) \cap T| = 2] \geq \Omega(1)$ whenever $x_{\{u,v\}} \not\approx \frac{1}{2}$. To do this, we split into two cases: when $x_{\{u,v\}} \geq \frac{1}{2} + \epsilon$ and when $x_{\{u,v\}} \leq \frac{1}{2} - \epsilon$.

▶ **Lemma 11.** *Let u, v be two vertices such that $x_{\{u,v\}} \geq \frac{1}{2} + \epsilon$ for some $\epsilon > 0$. Then,*

$$\mathbb{P}_{T\sim\mu}[|\delta(u) \cap T| = |\delta(v) \cap T| = 2] \geq \frac{\epsilon}{2e^3}.$$

Proof. Let $e = (u, v)$. For $T \sim \mu$, let $A_1 = \mathbb{I}[e \in T]$, $A_2 = |(\delta(u) \smallsetminus \{e\}) \cap T|$, $A_3 = |(\delta(v) \smallsetminus \{e\}) \cap T|$. We are now interested in the event $A_i = \kappa_i, \forall i$ for the vector $\kappa = (1, 1, 1)$ as this implies $|\delta(u) \cap T| = |\delta(v) \cap T| = 2$. We have $\mathbb{E}[A_1] = x_e$, $\mathbb{E}[A_2] = \mathbb{E}[A_3] = 2 - x_e$. Therefore, to apply Theorem 9 we can set:

$$\delta_1 = x_e, \delta_2 = 1, \delta_3 = 2 - x_e, \qquad \epsilon_1 = x_e, \epsilon_2 = 2x_e - 1, \epsilon_3 = x_e$$

In this case the worst case is using all of the ϵ terms in the bound, giving $e^{-3}x_e^2(2x_e - 1) \geq \frac{\epsilon}{2e^3}$ as desired. ◀

▶ **Lemma 12.** *Let u, v be two vertices such that $x_{\{u,v\}} \leq \frac{1}{2} - \epsilon$ for some $\epsilon > 0$. Then,*

$$\mathbb{P}_{T\sim\mu}[|\delta(u) \cap T| = |\delta(v) \cap T| = 2] \geq \frac{2\epsilon}{e^4}.$$

Proof. As above, $e = (u, v)$, and for $T \sim \mu$, let $A_1 = \mathbb{I}[e \in T]$, $A_2 = |(\delta(u) \smallsetminus \{e\}) \cap T|$, $A_3 = |(\delta(v) \smallsetminus \{e\}) \cap T|$. We are now interested in the event $A_i = \kappa_i, \forall i$ for the vector $\kappa = (0, 2, 2)$ as this implies $|\delta(u) \cap T| = |\delta(v) \cap T| = 2$. We have $\mathbb{E}[A_1] = x_e$, $\mathbb{E}[A_2] = \mathbb{E}[A_3] = 2 - x_e$. Therefore, to apply Theorem 9 we can set:

$$\delta_1 = 1 - x_e, \delta_2 = 1 - 2x_e, \delta_3 = 1 - x_e, \qquad \epsilon_1 = 1 - x_e, \epsilon_2 = 1, \epsilon_3 = 1 + x_e$$

In this case the worst case is using all of the δ terms in the bound, giving $4e^{-4}(1 - x_e)^2(1 - 2x_e) \geq \frac{2\epsilon}{e^4}$. ◀

[3] Technically, a spanning tree plus an edge is sampled, as otherwise one cannot exactly have $x \in P_{\text{Sub}}$, but we ignore that here.

We leave further background about TSP and the proofs of our improved probabilistic statements to the full version of the paper, as understanding their importance requires some knowledge of the (highly technical) proof in [24]. However, in the next section we summarize the new probabilistic bounds we get and their consequences.

4.2 Summary of Probabilistic Bounds and New Approximation Factor

The current bound on the performance of the max entropy algorithm is $\frac{3}{2} - 4.11 \cdot 10^{-36}$. This is primarily governed by a constant p that is determined by the minimum probability over a number of events. In [24], p was equal to $2 \cdot 10^{-10}$, and these events are described by the following statements in [24]:

1. Corollary 5.9, which gives a bound of $1.5 \cdot 10^{-9}$. We do not modify this bound, although we note that Lemma 5.7 can be slightly improved which would lead to a small improvement here.
2. Lemma 5.21, which gives a bound of $0.005\epsilon_{1/2}^2 = 2 \cdot 10^{-10}$. We improve this to $0.039\epsilon_{1/2}^2 = 1.56 \cdot 10^{-9}$.
3. Lemma 5.22, which gives a bound of $0.006\epsilon_{1/2}^2 = 2.4 \cdot 10^{-10}$. We improve this to $0.038\epsilon_{1/2}^2 = 1.52 \cdot 10^{-9}$.
4. Lemma 5.23, which gives a bound of $0.005 \cdot \epsilon_{1/2} = 2 \cdot 10^{-10}$. We observe here that arguments already in [24] can be used to give $0.0498\epsilon_{1/2}^2 \geq 1.9 \cdot 10^{-9}$.
5. Lemma 5.24, which gives a bound of $0.02\epsilon_{1/2}^2 = 8 \cdot 10^{-10}$. We improve this to $0.0485\epsilon_{1/2}^2 \geq 1.9 \cdot 10^{-9}$.
6. Lemma 5.27, which gives a bound of 0.01. This lemma actually uses that the threshold p is *small*, and therefore this bound decreases slightly upon raising p. However, as it is quite far from being the bottleneck in these bounds, we omit the proof that the probability remains above $1.5 \cdot 10^{-9}$.

Therefore, we may increase p to the minimum of all these probabilities, $1.5 \cdot 10^{-9}$. Using statements from [24, 23], the following then holds:

▶ **Lemma 13.** *Let p be a lower bound on the probabilities guaranteed by (1) - (6) for $\epsilon_{1/2} \leq 0.0002$, and suppose $p \leq 10^{-4}$. Then given $x \in P_{\text{Sub}}$, the max entropy algorithm returns a solution of expected cost at most $(\frac{3}{2} - 9.7p^2 \cdot 10^{-17}) \cdot c(x)$.*

As we improve the bounds on p to $1.5 \cdot 10^{-9}$, an immediate corollary is the following:

▶ **Corollary 14.** *The max entropy algorithm is a $\frac{3}{2} - 2.18 \cdot 10^{-34}$ approximation algorithm for metric TSP.*

Using [25] this guarantee can be made deterministic, as we do not require any modifications to the algorithm.

In the full version of the paper, we also observe a lower bound on Lemma 5.21 for strongly Rayleigh distributions of $\Omega(\epsilon_{1/2}^2)$. The fact that $\epsilon_{1/2} \leq 0.0002$ is used in many places in [24] and thus decreasing it may require more effort. Thus without modifying other parts of the argument, it may not be possible to improve the bound below $1.5 \cdot 10^{-31}$.

In the rest of the paper we prove our main capacity bound.

5 Proofs of the Main Capacity Bounds

In this section we prove the strongest forms of the main capacity bounds, which give Theorem 1 and Corollary 2 as corollaries. See the full version of the paper for any missing proofs. For this section, we utilize the following notation.

▶ **Definition 15.** *For $n, d \in \mathbb{N}$ and $\boldsymbol{\alpha} \in \mathbb{R}^n_{\geq 0}$, we define the following:*

1. $\mathrm{NHMat}^d_n(\boldsymbol{\alpha})$ *is the set of all $\mathbb{R}_{\geq 0}$-valued $d \times (n+1)$ matrices with row sums all equal to 1 and column sums equal to $\alpha_1, \ldots, \alpha_n, d - \|\boldsymbol{\alpha}\|_1$.*
2. $\mathrm{NHProd}^d_n(\boldsymbol{\alpha})$ *is the set of all polynomials of the form*

$$p(x) = \prod_{i=1}^{d} \left(a_{i,n+1} + \sum_{j=1}^{n} a_{i,j} x_j \right),$$

where $A \in \mathrm{NHMat}^d_n(\boldsymbol{\alpha})$. In this case, we call p the polynomial associated to A. *Note that $p(\mathbf{1}) = 1$ and $\nabla p(\mathbf{1}) = \boldsymbol{\alpha}$ for all such polynomials.*

3. $\mathrm{NHStab}^d_n(\boldsymbol{\alpha})$ *is the set of all real stable polynomials in $\mathbb{R}_{\geq 0}[x_1, \ldots, x_n]$ of degree at most d for which $p(\mathbf{1}) = 1$ and $\nabla p(\mathbf{1}) = \boldsymbol{\alpha}$. (Recall that a polynomial is stable if it is never zero when all inputs are in the open complex upper half-plane.)*

We also define the following for $n \in \mathbb{N}$, $\boldsymbol{\alpha} \in \mathbb{R}^n_{\geq 0}$, and $\boldsymbol{\kappa} \in \mathbb{Z}^n_{\geq 0}$:

$$L_n^{\mathrm{NHProd}}(\boldsymbol{\alpha}; \boldsymbol{\kappa}) := \min_{d \in \mathbb{N}} \min_{p \in \mathrm{NHProd}^d_n(\boldsymbol{\alpha})} \mathrm{Cap}_{\boldsymbol{\kappa}}(p).$$

We now follow the steps of the proof given in Section 3.

5.1 Productization for Non-homogeneous Stable Polynomials

We first show how we can reduce the problem of bounding the capacity to products of affine linear forms. We recall the main productization result from [22], which gives the non-homogeneous productization result as an immediate corollary.

▶ **Theorem 16** (Thm. 6.2, [22])**.** *Fix $n, d \in \mathbb{N}$, $\boldsymbol{u}, \boldsymbol{\alpha} \in \mathbb{R}^n_{\geq 0}$, and $p \in \mathbb{R}_{\geq 0}[x_1, \ldots, x_n]$ of homogeneous degree d, such that $p(\mathbf{1}) = 1$ and $\nabla p(\mathbf{1}) = \boldsymbol{\alpha}$. There exists an $\mathbb{R}_{\geq 0}$-valued $d \times n$ matrix A such that the rows sums of A are all equal to $\mathbf{1}$, the column sums of A are given by $\boldsymbol{\alpha}$, and $p(\boldsymbol{u}) = \prod_{i=1}^{d} (A\boldsymbol{u})_i$.*

▶ **Corollary 17.** *Fix $n, d \in \mathbb{N}$, $\boldsymbol{u}, \boldsymbol{\alpha} \in \mathbb{R}^n_{\geq 0}$, and $p \in \mathrm{NHStab}^d_n(\boldsymbol{\alpha})$. There exists $f \in \mathrm{NHProd}^d_n(\boldsymbol{\alpha})$ such that $p(\boldsymbol{u}) = f(\boldsymbol{u})$.*

Proof. Let $q(\boldsymbol{x}) = x_{n+1}^d \cdot p\left(\frac{x_1}{x_{n+1}}, \ldots, \frac{x_n}{x_{n+1}}\right)$ be the homogenization of p, and define $\boldsymbol{\beta} := \nabla q(\mathbf{1})$. So $q \in \mathbb{R}_{\geq 0}[x_1, \ldots, x_{n+1}]$ of homogeneous degree d such that $q(\mathbf{1}) = 1$ and $\nabla q(\mathbf{1}) = \boldsymbol{\beta} = (\alpha_1, \ldots, \alpha_n, d - \|\boldsymbol{\alpha}\|_1)$. Define $u_{n+1} := 1$, apply Theorem 16 to q and $\boldsymbol{u}$, and dehomogenize to obtain the desired result. ◀

We now use this result to reduce the problem of bounding the capacity to products of affine linear forms.

▶ **Corollary 18.** *For $p \in \mathrm{NHStab}^d_n(\boldsymbol{\alpha})$, we have*

$$\mathrm{Cap}_\kappa(p) \geq L_n^{\mathrm{NHProd}}(\boldsymbol{\alpha}; \boldsymbol{\kappa}).$$

Proof. For any $\boldsymbol{x} \in \mathbb{R}^n_{\geq 0}$, let $f \in \mathrm{NHProd}^d_n(\boldsymbol{\alpha})$ be such that $p(\boldsymbol{x}) = f(\boldsymbol{x})$ according to Corollary 17. With this, we have

$$\mathrm{Cap}_{\boldsymbol{\kappa}}(p) = \inf_{\boldsymbol{x} > 0} \frac{p(\boldsymbol{x})}{\boldsymbol{x}^{\boldsymbol{\kappa}}} \geq \inf_{\boldsymbol{x} > 0} \min_{d \in \mathbb{N}} \min_{f \in \mathrm{NHProd}^d_n(\boldsymbol{\alpha})} \frac{f(\boldsymbol{x})}{\boldsymbol{x}^{\boldsymbol{\kappa}}} = L_n^{\mathrm{NHProd}}(\boldsymbol{\alpha}; \boldsymbol{\kappa}).$$ ◀

5.2 The Extreme Points of $\mathrm{NHMat}_n^d(\boldsymbol{\alpha})$

The next result implies we can reduce to the extreme points of $\mathrm{NHMat}_n^d(\boldsymbol{\alpha})$ to lower bound $L_n^{\mathrm{NHProd}}(\boldsymbol{\alpha};\boldsymbol{\kappa})$.

▶ **Lemma 19** (See Thm. 3.1 of [10]). *Given $\boldsymbol{\kappa} \in \mathbb{Z}_{\geq 0}^n$ and $\boldsymbol{\alpha} \in \mathbb{R}_{\geq 0}^n$, let $\phi : \mathrm{NHMat}_n^d(\boldsymbol{\alpha}) \to \mathbb{R}_{\geq 0}$ be the function which maps M to $\mathrm{Cap}_{\boldsymbol{\kappa}}(p)$ where p is the polynomial associated to M. Then ϕ is log-concave on $\mathrm{NHMat}_n^d(\boldsymbol{\alpha})$.*

We next describe the extreme points of $\mathrm{NHMat}_n^d(\boldsymbol{\alpha})$ via bipartitioned forests.

▶ **Lemma 20.** *Any extreme point of $\mathrm{NHMat}_n^d(\boldsymbol{\alpha})$ has support given by a bipartite forest on d left vertices and $n+1$ right vertices.*

Proof. Let M be an extreme point of $\mathrm{NHMat}_n^d(\boldsymbol{\alpha})$, and suppose its bipartite support graph G does not give a forest. Then G must contain an even simple cycle. Group the edges of this cycle into two groups such that the odd edges make up one group, and the even edges make up the other (with any starting point). Add $\epsilon > 0$ to all matrix entries corresponding to even edges and subtract ϵ to all matrix entries corresponding to odd edges, to construct $M_+ \in \mathrm{NHMat}_n^d(\boldsymbol{\alpha})$. Do the same thing, but reverse the signs, to construct $M_- \in \mathrm{NHMat}_n^d(\boldsymbol{\alpha})$. Thus $M = \frac{M_+ + M_-}{2}$, contradicting the fact that M is an extreme point. ◀

In what follows, we will also need the following basic graph theoretic result.

▶ **Lemma 21.** *Let G be a bipartite forest on m left vertices and n right vertices such that G has no vertices of degree 0. Then G has at least $m - n + 1$ left leaves.*

5.3 Capacity Bounds via Induction

Here we complete the proof of Theorem 1. We first give a simple lemma, which bears some resemblance to the probabilistic union bound.

▶ **Lemma 22.** *Given $\boldsymbol{c} \in \mathbb{R}_{\geq 0}^d$ such that $c_i < 1$ for all $i \in [d]$, we have $1 - \sum_{i=1}^d c_i \leq \prod_{i=1}^d (1 - c_i)$.*

The next lemma handles the case from Step 4 in Section 3 of setting some variable equal to 0. To see how these next two lemmas actually are actually used, see Theorem 25 below.

▶ **Lemma 23.** *For $n \geq 1$, let $p \in \mathrm{NHProd}_n^d(\boldsymbol{\alpha})$ be the polynomial associated to a $d \times (n+1)$ matrix M, and suppose $\boldsymbol{\kappa} \in \mathbb{Z}_{\geq 0}^n$ such that $\kappa_n = 0$ and $\alpha_n - \kappa_n \leq 1 - \epsilon$ for some $\epsilon > 0$. Then there exists $q \in \mathrm{NHProd}_{n-1}^d(\boldsymbol{\beta})$ such that*

$$\mathrm{Cap}_{\boldsymbol{\kappa}}(p) \geq \epsilon \cdot \mathrm{Cap}_{\boldsymbol{\gamma}}(q),$$

where $\boldsymbol{\gamma} = (\kappa_1, \ldots, \kappa_{n-1}) \in \mathbb{Z}_{\geq 0}^{n-1}$ and $\boldsymbol{\beta} \in \mathbb{R}_{\geq 0}^{n-1}$ is such that for all $S \subseteq [n-1]$ we have

$$\sum_{j \in S} (\alpha_j - \kappa_j) \leq \sum_{j \in S} (\beta_j - \gamma_j) \leq (\alpha_n - \kappa_n) + \sum_{j \in S} (\alpha_j - \kappa_j).$$

The next lemma handles the case from Step 4 in Section 3 of taking the partial derivative with respect to some variable.

▶ **Lemma 24.** *For $n \geq 1$, let $p \in \mathrm{NHProd}_n^d(\boldsymbol{\alpha})$ be the polynomial associated to a $d \times (n+1)$ matrix M such that column n has exactly one non-zero entry, and suppose $\boldsymbol{\kappa} \in \mathbb{Z}_{\geq 0}^n$ is such that $\kappa_n = 1$ and $\kappa_n - \alpha_n \leq 1 - \epsilon$ for some $\epsilon > 0$. Then there exists $q \in \mathrm{NHProd}_{n-1}^{d-1}(\boldsymbol{\beta})$ such that*

$$\mathrm{Cap}_{\boldsymbol{\kappa}}(p) \geq \epsilon \cdot \mathrm{Cap}_{\boldsymbol{\gamma}}(q),$$

where $\boldsymbol{\gamma} = (\kappa_1, \ldots, \kappa_{n-1}) \in \mathbb{Z}_{\geq 0}^{n-1}$ and $\boldsymbol{\beta} \in \mathbb{R}_{\geq 0}^{n-1}$ is such that for all $S \subseteq [n-1]$ we have

$$\sum_{j \in S} (\kappa_j - \alpha_j) \leq \sum_{j \in S} (\gamma_j - \beta_j) \leq (\kappa_n - \alpha_n) + \sum_{j \in S} (\kappa_j - \alpha_j).$$

We now combine the lemmas above to prove our strongest capacity lower bound for polynomials in $\mathrm{NHProd}_n^d(\boldsymbol{\alpha})$.

▶ **Theorem 25.** *Fix any $\boldsymbol{\kappa} \in \mathbb{Z}_{\geq 0}^n$, $\boldsymbol{\alpha} \in \mathbb{R}_{\geq 0}^n$, and $\boldsymbol{\epsilon}, \boldsymbol{\delta} \in \mathbb{R}_{>0}^n$ such that*

$$\max_{S \in \binom{[n]}{k}} \sum_{j \in S} (\alpha_j - \kappa_j) \leq 1 - \epsilon_k \qquad \text{and} \qquad \max_{S \in \binom{[n]}{k}} \sum_{j \in S} (\kappa_j - \alpha_j) \leq 1 - \delta_k$$

for all $k \in [n]$. Then

$$\mathrm{Cap}_{\boldsymbol{\kappa}}(p) \geq \min_{0 \leq \ell \leq n} \prod_{k=1}^{\ell} \epsilon_k \prod_{k=1}^{n-\ell} \delta_k$$

for every $p \in \mathrm{NHProd}_n^d(\boldsymbol{\alpha})$.

The next result gives the bound which we will use in Section 4 to prove Theorem 6, the simply exponential improvement to the probability bound used for the metric TSP application.

▶ **Corollary 26.** *Fix any $\boldsymbol{\kappa} \in \mathbb{Z}_{\geq 0}^n$ and $\boldsymbol{\alpha} \in \mathbb{R}_{\geq 0}^n$ such that $\boldsymbol{\epsilon} \in \mathbb{R}_{>0}^n$ can be defined via*

$$\epsilon_k := 1 - \max_{\substack{S \subseteq [n] \\ |S| \leq k}} \left| \sum_{j \in S} (\kappa_j - \alpha_j) \right|$$

for all $k \in [n]$. Then $\mathrm{Cap}_{\boldsymbol{\kappa}}(p) \geq \prod_{k=1}^n \epsilon_k$ for every $p \in \mathrm{NHProd}_n^d(\boldsymbol{\alpha})$.

By Corollary 18, the above results hold for all $p \in \mathrm{NHStab}_n^d(\boldsymbol{\alpha})$, and we state this formally now.

▶ **Corollary 27.** *Theorem 25 and Corollary 26 hold for all $p \in \mathrm{NHStab}_n^d(\boldsymbol{\alpha})$.*

5.4 Proving Theorem 1 and Theorem 2

We now complete the proof of Theorem 1, and thus also of Corollary 2. Fix $p \in \mathrm{NHStab}_n^d(\boldsymbol{\alpha})$. Thus

$$1 - \max_{\substack{S \subseteq [n] \\ |S| \leq k}} \left| \sum_{j \in S} (\kappa_j - \alpha_j) \right| \geq 1 - \|\boldsymbol{\kappa} - \boldsymbol{\alpha}\|_1$$

and Corollary 27 (via Corollary 26) imply

$$\mathrm{Cap}_{\boldsymbol{\kappa}}(p) \geq (1 - \|\boldsymbol{\kappa} - \boldsymbol{\alpha}\|_1)^n,$$

which completes the proof. The tightness claim follows from Lemma 28.

5.5 Examples and Tightness

The δ parameters in Theorem 25 are tight, and this is shown in Example 29. However, the ϵ parameters are not, as shown in Example 30. Example 29 also demonstrates tightness of the dependence on the error parameter for some of our results, and we state this formally now.

▶ **Lemma 28.** *The dependence on $(1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)$ in Theorem 1 and Corollary 2 and the dependence on ϵ in Theorem 6 are all tight for any fixed $\boldsymbol{\kappa} > \mathbf{0}$.*

Proof. Define $\alpha_1 := \kappa_1 - (1-\epsilon) > 0$ and $\alpha_i := \kappa_i \geq 1$ for all $i \geq 2$. Let p be the polynomial described by Example 29, given explicitly by

$$p(\boldsymbol{x}) = \left(\prod_{i=1}^{n} x_i^{\kappa_i - 1}\right) \cdot \left(\prod_{i=1}^{n-1} (\epsilon x_i + (1-\epsilon) x_{i+1})\right) \cdot (\epsilon x_n + (1-\epsilon)).$$

Then $\boldsymbol{\kappa}$ is a vertex of of the Newton polytope of p. Thus

$$p_{\boldsymbol{\kappa}} = \mathrm{Cap}_{\boldsymbol{\kappa}}(p) = \prod_{k=1}^{n} \left(1 - \sum_{j=1}^{k} (\kappa_j - \alpha_j)\right) = \epsilon^n = (1 - \|\boldsymbol{\kappa} - \nabla p(\mathbf{1})\|_1)^n.$$

To see that p can be the probability generating polynomial for some random variables associated to a strongly Rayleigh distribution (as in Theorem 6), note that the polarization of p is real stable and gives the probability generating polynomial for such a strongly Rayleigh distribution. ◀

▶ **Example 29.** Fix any $\boldsymbol{\kappa} \in \mathbb{Z}_{\geq 0}^n$ and $\boldsymbol{\alpha} \in \mathbb{R}_{\geq 0}^n$ such that $\kappa_j - \alpha_j \geq 0$ for all $j \in [n]$ and $\|\boldsymbol{\kappa} - \boldsymbol{\alpha}\|_1 < 1$. Thus $\kappa_j \geq 1$ for all $j \in [n]$. For $d = \|\boldsymbol{\kappa}\|_1$, consider the matrix $M = \left[\begin{smallmatrix} A \\ B \end{smallmatrix}\right]$, where A is the $(\|\boldsymbol{\kappa}\| - n) \times (n+1)$ matrix given by

$$A = \begin{bmatrix} \mathbf{1}_{\kappa_1 - 1} \boldsymbol{e}_1^\top \\ \mathbf{1}_{\kappa_2 - 1} \boldsymbol{e}_2^\top \\ \vdots \\ \mathbf{1}_{\kappa_n - 1} \boldsymbol{e}_n^\top \end{bmatrix},$$

and B is the $n \times (n+1)$ matrix for which

$$b_{kk} = 1 - \sum_{j=1}^{k} (\kappa_j - \alpha_j), \qquad b_{k,k+1} = \sum_{j=1}^{k} (\kappa_j - \alpha_j),$$

for all $k \in [n]$ and $b_{ij} = 0$ otherwise. Note that every row sum of A is equal to 1, and the row sums of B are given by

$$\sum_{j=1}^{n} b_{kj} = 1 - \sum_{j=1}^{k} (\kappa_j - \alpha_j) + \sum_{j=1}^{k} (\kappa_j - \alpha_j) = 1$$

for all $k \in [n]$. The column sums of M are then given by

$$\sum_{i=1}^{d} m_{ik} = (\kappa_k - 1) + 1 - \sum_{j=1}^{k} (\kappa_j - \alpha_j) + \sum_{j=1}^{k-1} (\kappa_j - \alpha_j) = \kappa_k - (\kappa_k - \alpha_k) = \alpha_k$$

for all $k \in [n]$. Thus $M \in \mathrm{NHMat}_n^d(\boldsymbol{\alpha})$. Let p be the polynomial associated to M, and let q be the polynomial associated to B. We then have $\mathrm{Cap}_{\boldsymbol{\kappa}}(p) = \mathrm{Cap}_{\mathbf{1}}(q)$. Note that $\mathbf{1}$ is a vertex of the Newton polytope of q, and thus

$$\mathrm{Cap}_{\boldsymbol{\kappa}}(p) = \mathrm{Cap}_{\mathbf{1}}(q) = \prod_{k=1}^{n} b_{kk} = \prod_{k=1}^{n} \left(1 - \sum_{j=1}^{k} (\kappa_j - \alpha_j)\right).$$

By possibly permuting the variables to put $\kappa_j - \alpha_j$ in non-increasing order, this is precisely the lower bound guaranteed by Theorem 25.

▶ **Example 30.** Consider the case that $\boldsymbol{\kappa} = \mathbf{0}$ and $\|\boldsymbol{\alpha}\|_1 < 1$. Given any d, let M be any extreme point of $\mathrm{NHMat}_n^d(\boldsymbol{\alpha})$. Then the column sum of column $n+1$ of M is equal to $d - \sum_{j=1}^n \alpha_j > d - 1$. Since every row sum equals 1, all entries of column $n+1$ of M are strictly positive. Since M is an extreme point, this further implies that each column of M has at most 1 positive entry except column $n+1$. Letting p be the polynomial associated to M, there exists a partition $S_1 \sqcup \cdots \sqcup S_k = [n]$ such that

$$p(\boldsymbol{x}) = \prod_{i=1}^{k} \left(\left(\sum_{j \in S_i} \alpha_j x_j \right) + \left(1 - \sum_{j \in S_i} \alpha_j \right) \right).$$

Thus by Lemma 22,

$$\mathrm{Cap}_{\mathbf{0}}(p) = p_{\mathbf{0}} = \prod_{i=1}^{k} \left(1 - \sum_{j \in S_i} \alpha_j \right) \geq 1 - \sum_{j=1}^{n} \alpha_j.$$

By Lemma 19, this gives a lower bound on the capacity of every $p \in \mathrm{NHProd}_n^d(\boldsymbol{\alpha})$. However, this lower bound is strictly better than the one guaranteed by Theorem 25.

As a note, this can be partially remedied by removing all $\kappa_j = 0$ columns at the same time (i.e., adjusting Lemma 23 to remove many columns at once). However, it is currently unclear how to inductively do this correctly.

6 Uniqueness of Permanent Minimizers

In this section, let $\mathrm{Mat}_n(\boldsymbol{c})$ be the set of $n \times n$ matrices with non-negative entries and rows sums $\mathbf{1}$ and column sums $\boldsymbol{c} > 0$, let $p_M(\boldsymbol{x}) := \prod_{i=1}^n \sum_{j=1}^n m_{ij} x_j$ be the real stable polynomial associated to a given $M \in \mathrm{Mat}_n(\boldsymbol{c})$, and let $L(\boldsymbol{c})$ be any lower bound on $\mathrm{Cap}_{\mathbf{1}}(p_M)$ over all $M \in \mathrm{Mat}_n(\boldsymbol{c})$ (e.g., as given by Theorem 1 above or the results of [22]).

We now prove Theorem 3 via Theorem 32 below. We note that the argument in the proof of Theorem 32 given below can be made into a general statement about minimizers of quadratic forms. And further, the same argument given here applies to mixed discriminants, as mentioned in Section 2.

▶ **Lemma 31.** *If* $\frac{(n-2)^{n-2} n^{n-1}}{(n-1)^{2n-3}} > \frac{c_1 c_2 \cdots c_n}{L(\boldsymbol{c})}$, *then all minimizers of the permanent over over* $\mathrm{Mat}_n(\boldsymbol{c})$ *have all strictly positive entries.*

Proof. Note that the rank-one matrix $\frac{1}{n}\mathbf{1} \cdot \boldsymbol{c}^\top$ has all strictly positive entries and permanent equal to $\frac{n!}{n^n} \prod_{i=1}^n c_i$. Thus, to obtain a contradiction, let us assume that there exists $M \in \mathrm{Mat}_n(\boldsymbol{c})$ with at least one zero entry such that $\mathrm{per}(M) \leq \frac{n!}{n^n} \prod_{i=1}^n c_i$. By the main result of [18], we have

$$\frac{n!}{n^n} \prod_{i=1}^{n} c_i \geq \mathrm{per}(M) \geq \left(\frac{n-2}{n-1} \right)^{n-2} \frac{(n-1)!}{(n-1)^{n-1}} \cdot L(\boldsymbol{c}),$$

which after rearranging implies

$$\frac{(n-2)^{n-2}n^{n-1}}{(n-1)^{2n-3}} > \left(\frac{n-2}{n-1}\right)^{n-2} \frac{n^{n-1}}{(n-1)^{n-1}},$$

which is a contradiction. ◀

▶ **Theorem 32.** *If $\frac{(n-2)^{n-2}n^{n-1}}{(n-1)^{2n-3}} > \frac{c_1 c_2 \cdots c_n}{L(\boldsymbol{c})}$, then the rank-one matrix $\frac{1}{n}\mathbf{1} \cdot \boldsymbol{c}^\top$ is the unique minimizer of the permanent over* $\mathrm{Mat}_n(\boldsymbol{c})$.

Proof. Let M be a minimizer of the permanent over $\mathrm{Mat}_n(\boldsymbol{c})$. Thus every entry of M is positive by Lemma 31. We will show that every pair of rows of M must be equal, which immediately implies $M = \frac{1}{n}\mathbf{1} \cdot \boldsymbol{c}^\top$.

Let $P \subset \mathbb{R}^{2\times n}$ be the convex polytope given by the first two rows of all matrices in $\mathrm{Mat}_n(\boldsymbol{c})$ with rows 3 through n equal to those of M. Thus by positivity, $(\boldsymbol{m}_1, \boldsymbol{m}_2)$ is in the relative interior of P where $\boldsymbol{m}_i$ is the i^{th} row of M. Let $f(\boldsymbol{u}, \boldsymbol{v})$ be defined as the permanent of the matrix M with the first two rows replaced by $(\boldsymbol{u}, \boldsymbol{v})$. Since M minimizes the permanent and $(\boldsymbol{m}_1, \boldsymbol{m}_2)$ is contained in the relative interior of P, the necessary conditions on the minimum given by Lagrange multipliers implies

$$\nabla f(\boldsymbol{m}_1, \boldsymbol{m}_2) = \big(\mathrm{per}(M_{(i,j)})\big)_{i\in[2], j\in[n]} = (a_i + b_j)_{i\in[2], j\in[n]}$$

for some a_i, b_j, where $M_{(i,j)}$ is the matrix M with row i and column j deleted. That is, the gradient of f points in a direction orthogonal to P at $(\boldsymbol{m}_1, \boldsymbol{m}_2)$.

By row symmetry of the permanent, we also have that $\nabla f(\boldsymbol{m}_2, \boldsymbol{m}_1) = (a_i' + b_j)_{i\in[2], j\in[n]}$, where $a_1' = a_2$ and $a_2' = a_1$. By row multilinearity of the permanent, we thus have

$$\nabla f(t \cdot \boldsymbol{m}_1 + (1-t) \cdot \boldsymbol{m}_2, (1-t) \cdot \boldsymbol{m}_1 + t \cdot \boldsymbol{m}_2) = (t \cdot a_i + (1-t) \cdot a_i' + b_j)_{i\in[2], j\in[n]}$$

for all $t \in \mathbb{R}$. If $\boldsymbol{m}_1 \neq \boldsymbol{m}_2$ then the gradient of f is orthogonal to P at all points on a line in P through M, and thus the permanent is minimized at all these points. At least one such point (on the boundary of P) has a zero entry, which contradicts Lemma 31. Therefore it must be that $\boldsymbol{m}_1 = \boldsymbol{m}_2$. Applying this argument to every pair of rows of M implies the desired result. ◀

We now give the example which proves Proposition 4.

▶ **Example 33.** Fix $t > 0$ and $n \in \mathbb{N}$, and define $\epsilon := \frac{1}{n^{1+t}}$. Further define $\boldsymbol{\alpha} \in \mathbb{R}^n_{>0}$ and $\boldsymbol{c} \in \mathbb{R}^{n+1}_{>0}$ via

$$\boldsymbol{\alpha} := \big(1-\epsilon, 1-\epsilon, \ldots, 1-\epsilon\big) \in \mathbb{R}^n_{>0} \quad \text{and} \quad \boldsymbol{c} := \Big(1+\alpha_1, \alpha_2, \alpha_3, \ldots, \alpha_n, \sum_{j=1}^n (1-\alpha_j)\Big) \in \mathbb{R}^{n+1}_{>0}.$$

Note that $\|\boldsymbol{c} - \mathbf{1}\|_1 = 1 + (n-2)\epsilon + 1 - n\epsilon = 2(1-\epsilon) < 2$. We first have

$$\mathrm{per}\left(\frac{1}{n+1}\mathbf{1} \cdot \boldsymbol{c}^\top\right) = \frac{(n+1)!}{(n+1)^{n+1}}(2-\epsilon)(1-\epsilon)^{n-1}n\epsilon = \frac{n!}{n^{(1+t)n}} \cdot \frac{2-n^{-1-t}}{1-n^{-1-t}} \cdot \frac{(n^{1+t}-1)^n}{n^t(n+1)^n}.$$

Now consider the $(n+1) \times (n+1)$ matrix with diagonal entries

$$1, \quad \sum_{j=1}^{1}(1-\alpha_j), \quad \sum_{j=1}^{2}(1-\alpha_j), \quad \cdots, \quad \sum_{j=1}^{n}(1-\alpha_j),$$

subdiagonal entries to obtain row sums $\mathbf{1}$ and column sums $\boldsymbol{c}$, and zero entries elsewhere. Note that $M \in \mathrm{Mat}_{n+1}(\boldsymbol{c})$. The matrix M is upper-triangular, and thus we have

$$\mathrm{per}(M) = \prod_{k=1}^{n} \sum_{j=1}^{k} (1 - \alpha_j) = \prod_{k=1}^{n} (k\epsilon) = n! \cdot \epsilon^n = \frac{n!}{n^{(1+t)n}}.$$

We then further have

$$\frac{2 - n^{-1-t}}{1 - n^{-1-t}} \cdot \frac{(n^{1+t} - 1)^n}{n^t (n+1)^n} \approx 2n^{(n-1)t} \cdot \frac{(n - n^{-t})^n}{(n+1)^n} \geq 2n^{(n-1)t} \cdot \frac{(n-1)^n}{(n+1)^n} > 1,$$

which implies $\mathrm{per}(M) < \mathrm{per}\left(\frac{1}{n+1}\mathbf{1} \cdot \boldsymbol{c}^\top\right)$ for large enough n.

References

1 Kareem Adiprasito, June Huh, and Eric Katz. Hodge theory for combinatorial geometries. *Annals of Mathematics*, 188(2):381–452, 2018.

2 Yeganeh Alimohammadi, Nima Anari, Kirankumar Shiragur, and Thuy-Duong Vuong. Fractionally log-concave and sector-stable polynomials: counting planar matchings and more. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 433–446, 2021.

3 Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials iii: Mason's ultra-log-concavity conjecture for independent sets of matroids. *arXiv preprint*, 2018. `arXiv:1811.01600`.

4 Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials ii: High-dimensional walks and an fpras for counting bases of a matroid. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 1–12, 2019.

5 Nima Anari, Tung Mai, Shayan Oveis Gharan, and Vijay V Vazirani. Nash social welfare for indivisible items under separable, piecewise-linear concave utilities. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2274–2290. SIAM, 2018.

6 Nima Anari and Shayan Oveis Gharan. The kadison-singer problem for strongly rayleigh measures and applications to asymmetric tsp. *arXiv preprint*, 2014. `arXiv:1412.1143`.

7 Nima Anari and Shayan Oveis Gharan. A generalization of permanent inequalities and applications in counting and optimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 384–396, 2017.

8 Nima Anari, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials, entropy, and a deterministic approximation algorithm for counting bases of matroids. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 35–46. IEEE, 2018.

9 Arash Asadpour, Michel X Goemans, Aleksander Madry, Shayan Oveis Gharan, and Amin Saberi. An o (log n/log log n)-approximation algorithm for the asymmetric traveling salesman problem. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete algorithms*, pages 379–389, 2010.

10 Alexander Barvinok. Enumerating contingency tables via random permanents. *Combinatorics, Probability and Computing*, 17(1):1–19, 2008.

11 Alexander Barvinok. Computing the permanent of (some) complex matrices. *Foundations of Computational Mathematics*, 16:329–342, 2016.

12 Julius Borcea, Petter Brändén, and Thomas Liggett. Negative dependence and the geometry of polynomials. *Journal of the American Mathematical Society*, 22(2):521–567, 2009.

13 Petter Brändén. Hyperbolic polynomials and the kadison-singer problem. *arXiv preprint*, 2018. `arXiv:1809.03255`.

14 Petter Brändén and June Huh. Lorentzian polynomials. *Annals of Mathematics*, 192(3):821–891, 2020.

15 Petter Brändén, Jonathan Leake, and Igor Pak. Lower bounds for contingency tables via lorentzian polynomials. *Israel Journal of Mathematics*, 253(1):43–90, 2023.

16 Gregory P Egorychev. The solution of van der waerden's problem for permanents. *Advances in Mathematics*, 42(3):299–305, 1981.

17 Dmitry I Falikman. Proof of the van der waerden conjecture regarding the permanent of a doubly stochastic matrix. *Mathematical notes of the Academy of Sciences of the USSR*, 29:475–479, 1981.

18 Leonid Gurvits. Hyperbolic polynomials approach to van der waerden/schrijver-valiant like conjectures: sharper bounds, simpler proofs and algorithmic applications. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 417–426, 2006.

19 Leonid Gurvits. The van der waerden conjecture for mixed discriminants. *Advances in Mathematics*, 200(2):435–454, 2006.

20 Leonid Gurvits. On multivariate newton-like inequalities. In *Advances in Combinatorial Mathematics: Proceedings of the Waterloo Workshop in Computer Algebra 2008*, pages 61–78. Springer, 2009.

21 Leonid Gurvits. A polynomial-time algorithm to approximate the mixed volume within a simply exponential factor. *Discrete & Computational Geometry*, 41:533–555, 2009.

22 Leonid Gurvits and Jonathan Leake. Capacity lower bounds via productization. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 847–858, 2021.

23 Anna Karlin, Nathan Klein, and Shayan Oveis Gharan. A (slightly) improved bound on the integrality gap of the subtour lp for tsp. In *FOCS*, pages 844–855. IEEE Computer Society, 2022.

24 Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. A (slightly) improved approximation algorithm for metric tsp. In *STOC*. ACM, 2021.

25 Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. A deterministic better-than-3/2 approximation algorithm for metric tsp. In Alberto Del Pia and Volker Kaibel, editors, *Integer Programming and Combinatorial Optimization*, pages 261–274, Cham, 2023. Springer International Publishing.

26 Adam W Marcus, Daniel A Spielman, and Nikhil Srivastava. Interlacing families ii: Mixed characteristic polynomials and the kadison—singer problem. *Annals of Mathematics*, pages 327–350, 2015.

27 Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. A randomized rounding approach to the traveling salesman problem. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 550–559. IEEE, 2011.

28 Damian Straszak and Nisheeth K Vishnoi. Real stable polynomials and matroids: Optimization and counting. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 370–383, 2017.