

A Tight Analysis of Bethe Approximation for Permanent

Nima Anari
 Computer Science Department
 Stanford University
 anari@cs.stanford.edu

Alireza Rezaei
 Computer Science & Engineering
 University of Washington
 arezaei@cs.washington.edu

Abstract—We prove that the permanent of nonnegative matrices can be deterministically approximated within a factor of $\sqrt{2}^n$ in polynomial time, improving upon the previous deterministic approximations. We show this by proving that the Bethe approximation of the permanent, a quantity computable in polynomial time, is at least as large as the permanent divided by $\sqrt{2}^n$. This resolves a conjecture of Gurvits [21]. Our bound is tight, and when combined with previously known inequalities lower bounding the permanent, fully resolves the quality of Bethe approximation for permanent.

Keywords-permanent; deterministic counting; bethe approximation;

I. INTRODUCTION

The *permanent* of a square matrix has been an important object of study in combinatorics and theoretical computer science. The permanent of an $n \times n$ matrix A is defined as

$$\text{per}(A) = \sum_{\sigma \in S_n} \prod_{i=1}^n A_{i,\sigma(i)},$$

where S_n is the set of all permutations $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$. Although this definition of the permanent looks superficially similar to that of the determinant, the latter is polynomial-time computable while the former is $\#P$ -hard to compute even when A has entries in $\{0, 1\}$ [33].

Given the hardness of computing the permanent exactly, a natural question is when can the permanent be approximated efficiently? For arbitrary $A \in \mathbb{C}^{n \times n}$, an $\epsilon \|A\|^n$ additive approximation¹ was given by Gurvits [20]; this was later improved and derandomized in certain cases [2, 12]. The permanent of complex matrices is a central quantity in quantum linear optics, and is connected to some methods proposed for demonstrating “quantum supremacy” [1].

A *multiplicative* approximation is elusive for general A ; even determining the sign of the permanent of real matrices is hard [1, 19]. This motivates studying classes of matrices where the sign is not a barrier. Polynomial time multiplicative approximation algorithms have been designed for two large classes known to have a nonnegative permanent: the most prominent is a randomized $(1 + \epsilon)$ -approximation algorithm for matrices with nonnegative entries [24], and there is also a $\simeq 4.84^n$ -approximation algorithm for positive

semidefinite matrices [4]. We remark that there are cases where the permanent is not even necessarily real-valued, yet quasi-polynomial time multiplicative $(1 + \epsilon)$ -approximation has been achieved [9, 8, 18]. In this paper we focus only on matrices with nonnegative entries; for these, permanent has a close connection to the combinatorics of perfect matchings in bipartite graphs.

For $A \in \{0, 1\}^{n \times n}$, the permanent of A is equal to the number of perfect matchings in the bipartite graph with bipartite adjacency matrix A . For nonnegative $A \in \mathbb{R}_{\geq 0}^{n \times n}$ the permanent is clearly a nonnegative quantity, and in fact can be thought of as a partition function. Consider sampling $\sigma \in S_n$ with probability $\mathbb{P}[\sigma] \propto \prod_{i=1}^n A_{i,\sigma(i)}$. Then $\text{per}(A)$ is the normalizing constant in this model, i.e., $\mathbb{P}[\sigma] = (\prod_{i=1}^n A_{i,\sigma(i)}) / \text{per}(A)$. When $A \in \{0, 1\}^{n \times n}$, this distribution over $\sigma \in S_n$ is the same as the uniform distribution over perfect matchings in the corresponding bipartite graph.

Given the close connection between random perfect matchings and the permanent of nonnegative matrices, Broder [11] suggested using Monte Carlo Markov Chain (MCMC) methods to approximate the permanent. This was finally achieved by Jerrum, Sinclair, and Vigoda [24], who gave a fully polynomial time randomized approximation scheme (FPRAS); that is a randomized algorithm that for any desired $\epsilon, \delta > 0$, with probability $1 - \delta$ computes a $1 + \epsilon$ approximation to $\text{per}(A)$ for nonnegative A in time $\text{poly}(n, \langle A \rangle, 1/\epsilon, \log(1/\delta))$, where $\langle A \rangle$ represents the bit complexity of A .

A. Deterministic approximation of the permanent

Despite having a complete understanding of approximations achievable by randomized algorithms, there is currently a large gap in our understanding of deterministic permanent approximation. The only known lower bound is still the $\#P$ -hardness of exact computation [33]. On the other hand, the best known deterministic approximation ratio is $\simeq 1.9022^n \leq 2^n$ [22].

Deterministic approximations of the permanent can sometimes be useful in optimization tasks. For example, Anari et al. [6] used a particular approximation of the permanent to attack the problem of maximizing Nash social welfare; a key property used in their algorithm was that this approximation

¹We remark that for any matrix A , $|\text{per}(A)| \leq \|A\|^n$.

of $\text{per}(A)$ was log-concave as a function of the matrix A , so it could be maximized for A ranging over a convex set of matrices. Another problem that can be cast as permanent maximization is that of finding the profile maximum likelihood, a successful universal method for estimating symmetric statistics from sampled data [see 3]. One of the successful heuristics used to find the profile maximum likelihood estimate is based on a particular approximation of the permanent called the Bethe approximation [34]. In fact this is the same approximation we study in this paper. For more recent results on profile maximum likelihood see [3, 28, 13].

It is no coincidence that the gap between randomized and deterministic approximations of permanent is so large: $1 + \epsilon$ vs. exponential. This is because, roughly speaking, no approximation ratio can be in between the two extremes.² More precisely, any polynomial time 2^{n^c} -approximation for any constant $c < 1$ can be boosted to obtain a $(1 + \epsilon)$ -approximation. This is achieved by feeding the hypothetical algorithm the larger matrix

$$I_m \otimes A = \begin{bmatrix} A & 0 & \dots & 0 \\ 0 & A & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & A \end{bmatrix},$$

whose size is $mn \times mn$ and whose permanent is $\text{per}(A)^m$. By taking the m -th root of the hypothetical algorithm's approximation of $\text{per}(I_m \otimes A)$, we obtain a $\sqrt[m]{2^{(mn)^c}}$ -approximation to $\text{per}(A)$. It is enough to take $m \simeq (n^c/\epsilon)^{1/(1-c)} = \text{poly}(n, 1/\epsilon)$ for this approximation ratio to become $1 + \epsilon$.

This suggests that the right question to ask about the approximation ratio is: What is the infimum of constants $\alpha > 1$ for which the permanent can be deterministically approximated within a factor of α^n in polynomial time? Linial, Samorodnitsky, and Wigderson [25] achieved the first nontrivial deterministic approximation of the permanent and showed that $\alpha \leq e$. Later Gurvits and Samorodnitsky [22] showed³ that $\alpha \leq 1.9022 < 2$. We prove that $\alpha \leq \sqrt{2}$.

Theorem 1. *There is a deterministic polynomial time algorithm that approximates the permanent of nonnegative matrices within a factor of $\sqrt{2}^n$.*

B. Bethe approximation of the permanent

The algorithm we analyze has, in fact, been studied very extensively and its output is known as the Bethe approximation of the permanent, or Bethe permanent for short [35].

²We remark that this does not rule out the possibility of the best approximation ratio being for example $2^{n/\text{poly} \log(n)}$.

³The approximation ratio obtainable by the analysis of Gurvits and Samorodnitsky [22] is actually $\simeq 1.9022^n$ but this is not explicitly mentioned in their article, and the only explicit bound mentioned is 2^n .

Definition 2. For a matrix $A \in \mathbb{R}_{\geq 0}^{n \times n}$, define the Bethe permanent of A as

$$\text{bethe}(A) = \max \left\{ \prod_{i,j=1}^n (A_{i,j}/P_{i,j})^{P_{i,j}} (1 - P_{i,j})^{1-P_{i,j}} \mid \begin{array}{l} P \in \mathbb{R}_{\geq 0}^{n \times n}, \\ P\mathbf{1} = \mathbf{1}, \\ \mathbf{1}^\top P = \mathbf{1}^\top. \end{array} \right\}$$

This maximization is defined over *doubly stochastic* matrices P ; that is nonnegative matrices whose rows and columns sum to 1.

Bethe approximation is a heuristic used to approximate the partition function of arbitrary graphical models, whose roots go back to statistical physics. Bethe approximation is very closely connected to the sum-product algorithm, also known as belief propagation [36, 35].

Natural graphical models can be easily constructed where the permanent is the partition function and the corresponding Bethe approximation, the Bethe permanent. It was observed that the sum-product algorithm on these graphical models yields a good approximation to the permanent in practice [15, 23]. Later Vontobel [35] showed that for these graphical models the sum-product algorithm always converges to the maximizer in definition 2, and that the maximization objective in definition 2 becomes a concave function of P after taking log. So besides the sum-product algorithm, $\text{bethe}(A)$ can also be computed by convex programming methods. We remark that this is one of the rare instances where the sum-product algorithm is *guaranteed* to converge.

The first theoretical guarantee for the quality of Bethe approximation was actually obtained by Gurvits and Samorodnitsky [22]. Their 2^n -approximation guarantee was for a slightly modified variant of the Bethe permanent. In particular, Gurvits [21] showed that $\text{bethe}(A)$ always lower bounds $\text{per}(A)$ and later Gurvits and Samorodnitsky [22] showed that $2^n \text{bethe}(A)$ upper bounds $\text{per}(A)$. The lower bound was proved by using an inequality of Schrijver [31]. Two alternative proofs of the lower bound have since been obtained based on 2-lifts [16], and real stable polynomials [5]. For a survey of these proofs and their generalizations beyond permanent see Straszak and Vishnoi [32].

Theorem 3 ([21] based on [31]). *For any matrix $A \in \mathbb{R}_{\geq 0}^{n \times n}$ we have*

$$\text{per}(A) \geq \text{bethe}(A).$$

While the lower bound is an elegant inequality that is also tight [22], the upper bound of $2^n \text{bethe}(A)$ had room for improvement. In fact Gurvits [21] had already conjectured that the correct constant for the upper bound should be $\sqrt{2}^n$. We resolve this conjecture.

Theorem 4. *For any matrix $A \in \mathbb{R}_{\geq 0}^{n \times n}$ we have*

$$\text{per}(A) \leq \sqrt{2}^n \text{bethe}(A).$$

Theorems 3 and 4 together with polynomial time computability [35, 22] imply theorem 1.

C. Techniques for upper bounding permanent

Perhaps the most well-known upper bound on the permanent is the Bregman-Minc inequality, conjectured by Minc [27] and proved by Bregman [10], which states that for $A \in \{0, 1\}^{n \times n}$ if the row sums are d_1, \dots, d_n , then

$$\text{per}(A) \leq \sqrt[n]{d_1! \cdots d_n!}.$$

This bound has been reproved by Schrijver [30] and Radhakrishnan [29] using entropy-based methods. In these works, some generalizations have been suggested for non-binary matrices $A \in \mathbb{R}_{\geq 0}^{n \times n}$, but as far as we know none imply theorem 4. We remark that entropy-based bounds have recently found applications in counting bases of matroids and matroid intersections [7]; perfect matchings are a special case of matroid intersections.

Our method for proving theorem 4 is inspired by the entropy-based proofs of the Bregman-Minc inequality, but differs from them in key places. To illustrate, prior entropy-based proofs of Bregman-Minc proceeded by comparing, in terms of entropy-based measures, two distributions on perfect matchings in the bipartite graph associated with A . One was the uniform distribution μ over perfect matchings, and the other one ν can be described by the following sampling procedure: Iterate over vertices on one side, and each time pick a neighbor for the current vertex uniformly at random out of all currently unmatched neighbors. We remark that the latter distribution has been used to estimate the permanent by importance sampling, and an entropy-based measure of distance between μ and ν has been shown to have a close connection to the sample complexity [17, 14].

The generalization of μ to arbitrary $A \in \mathbb{R}_{\geq 0}^{n \times n}$ is straightforward. We just let $\mu(\sigma) \propto \prod_{i=1}^n A_{i, \sigma(i)}$ for $\sigma \in S_n$. Similarly there is a natural generalization of ν : In iteration i we just sample an unmatched neighbor j with probability proportional to $A_{i, j}$. Indeed this generalization has been studied by Schrijver [30] in order to obtain generalizations of Bregman-Minc to non-binary matrices.

Our main insight is that defining ν according to another matrix $P \in \mathbb{R}_{\geq 0}^{n \times n}$ can result in a distribution that is closer to μ . To be more precise, we find a matrix P other than A , and modify the sampling procedure defining ν to pick a neighbor j for i with probability $\propto P_{i, j}$. We will ultimately let P be the marginals of μ , i.e., $P_{i, j} = \mathbb{P}_{\sigma \sim \mu}[\sigma(i) = j]$. One feature of this modification for example is that the modified ν never picks any edge that can never be part of a sample of μ . Comparing μ and the modified ν ultimately allows us to prove theorem 4.

Our technique for upper bounding the permanent is completely analytical and is described in section IV. However the approximation ratio guarantee resulting from it is only described in terms of the maximum of a well-defined

universal function on the standard simplex. The maxima of this function can be easily guessed; but in order to prove that they are indeed the maxima, further ideas are needed. We combine several analytic ideas to reduce the possible set of maxima to a finite set (of moderate size) and with the help of a computer search formally prove that the guessed maxima of this function are indeed the global maxima, yielding the desired approximation ratio of $\sqrt{2}^n$. We remark that even though we used a computer search, we analytically show that the search can be done over a particular finite set of points and with exact rational arithmetic. So the computer program's successful execution should be considered a formal *certificate* of correctness. For details of the computer search and the proof that the search can be reduced to a finite set see section V.

II. NOTATION AND PRELIMINARIES

Everywhere \log is taken in base e . We use $[n]$ to denote the set $\{1, \dots, n\}$. For a set S we denote the family of all subsets of S by 2^S . We use the notation $\mathbf{1}$ to denote the vector of all ones in \mathbb{R}^n . For a set $S \subseteq [n]$, we use $\mathbf{1}_S$ to denote the indicator vector of S which belongs to $\{0, 1\}^n$. We use Δ_n to denote the standard simplex⁴ of all n -dimensional stochastic vectors, i.e., all $p \in \mathbb{R}_{\geq 0}^n$ such that $\mathbf{1}^\top p = 1$:

$$\Delta_n = \{(p_1, \dots, p_n) \in \mathbb{R}_{\geq 0}^n \mid p_1 + \dots + p_n = 1\}.$$

We call a matrix $P \in \mathbb{R}_{\geq 0}^{n \times n}$ doubly stochastic when each row and each column of P belongs to Δ_n , i.e., $P\mathbf{1} = \mathbf{1}$ and $\mathbf{1}^\top P = \mathbf{1}^\top$. The set of all doubly stochastic matrices is known as the Birkhoff polytope which we denote by \mathcal{B}_n :

$$\mathcal{B}_n = \{P \in \mathbb{R}_{\geq 0}^{n \times n} \mid P\mathbf{1} = \mathbf{1} \wedge \mathbf{1}^\top P = \mathbf{1}^\top\}.$$

The fact that \mathcal{B}_n is a polytope can be easily seen by noticing that it is defined by $2n$ linear equalities and n^2 linear inequalities. An alternative description of \mathcal{B}_n is as the convex hull of permutation matrices [see, e.g., 37].

For a pair $e = (i, j) \in [n] \times [n]$, we use $A_e = A_{i, j}$ to denote the entry in the i -th row and j -th column of A . We use S_n to denote the set of all permutations $\sigma : [n] \rightarrow [n]$. With some abuse of notation we also view σ as the set $\{(i, \sigma(i)) \mid i \in [n]\} \subseteq [n] \times [n]$. So in particular $e \in \sigma$ means that e is a pair of the form $(i, \sigma(i))$ for some $i \in [n]$. With this notation the permanent can be written as

$$\text{per}(A) = \sum_{\sigma \in S_n} \prod_{e \in \sigma} A_e.$$

We sometimes think of $\sigma \in S_n$ as a perfect matching in a bipartite graph: The first side is the domain of σ , and the other side the image of σ . For two permutations $\sigma, \pi \in S_n$ we take $\sigma \circ \pi$ to be the composed permutation defined by $\sigma \circ \pi$

⁴We remark that in geometry this set is usually denoted Δ_{n-1} but we deviate from the standard slightly and use Δ_n .

$\pi(i) = \sigma(\pi(i))$. We also sometimes view a permutation $\pi \in S_n$ as a total ordering on $\{1, \dots, n\}$, that is we define the ordering $<_\pi$ on $[n]$ by $\pi(1) <_\pi \pi(2) <_\pi \dots <_\pi \pi(n)$. For a permutation π , we let $\bar{\pi}$ denote the reverted permutation defined by $\bar{\pi}(i) = \pi(n - i + 1)$. The total ordering defined by π and $\bar{\pi}$ are the opposites of each other.

For two distributions μ and ν supported on a finite set Ω , we define their Kullback-Leibler divergence, KL divergence for short, as

$$\mathcal{D}(\mu \parallel \nu) = \sum_{\omega \in \Omega} \mu(\omega) \log \frac{\mu(\omega)}{\nu(\omega)} = \mathbb{E}_{\omega \sim \mu} \left[\log \frac{\mu(\omega)}{\nu(\omega)} \right].$$

An important well-known property of the KL divergence is nonnegativity. This can be proved for example by applying Jensen's inequality to the convex function $\log(1/\cdot)$ [see, e.g., 26].

Fact 5 (Gibbs' inequality). *For any two distributions μ and ν we have $\mathcal{D}(\mu \parallel \nu) \geq 0$.*

III. BETHE PERMANENT

Doubly stochastic matrices have an intimate connection to random permutations. They can be thought of as marginals of distributions on permutations. More precisely, for any distribution μ on S_n , the $n \times n$ matrix P defined by $P_{i,j} = \mathbb{P}_{\sigma \sim \mu}[(i, j) \in \sigma]$ is doubly stochastic. Alternatively for any $P \in \mathcal{B}_n$, there is μ supported on S_n whose marginals are P ; this is because the vertices of \mathcal{B}_n are permutation matrices.

Given a matrix $A \in \mathbb{R}_{\geq 0}^{n \times n}$, a natural distribution μ on S_n can be defined by $\mu(\sigma) \propto \prod_{e \in \sigma} A_e$. The partition function, or the normalizing constant in this definition is $\text{per}(A)$. Given these connections, it is no surprise that many upper and lower bounds on the $\text{per}(A)$ can be expressed in terms of doubly stochastic matrices. For example the upper bound implicitly used by Linial, Samorodnitsky, and Wigderson [25] is

$$\log \text{per}(A) \leq \max \left\{ \sum_{e \in [n] \times [n]} P_e \log(A_e/P_e) \mid P \in \mathcal{B}_n \right\},$$

which can be interpreted as a form of subadditivity of entropy. The Van der Waerden conjecture used as a lower bound by Linial, Samorodnitsky, and Wigderson [25] can be rewritten as

$$n + \log \text{per}(A) \geq \sum_{e \in [n] \times [n]} P_e \log(A_e/P_e),$$

for all $P \in \mathcal{B}_n$. For a more detailed account of these connections see Gurvits [21]. We remark that the P maximizing the r.h.s. of the above can also be obtained by matrix scaling. In other words, as long as $\text{per}(A) > 0$, there are sequences

of diagonal matrices R_1, R_2, \dots and C_1, C_2, \dots such that

$$\lim_{i \rightarrow \infty} R_i A C_i = \text{argmax} \left\{ \sum_{e \in [n] \times [n]} P_e \log(A_e/P_e) \mid P \in \mathcal{B}_n \right\}.$$

Chertkov, Kroc, and Vergassola [15] and Huang and Jebara [23] observed that the μ distribution defined for a matrix A can be described by a graphical model. Roughly speaking, one can define factors $R_1, \dots, R_n : 2^{[n] \times [n]} \rightarrow \mathbb{R}_{\geq 0}$ and $C_1, \dots, C_n : 2^{[n] \times [n]} \rightarrow \mathbb{R}_{\geq 0}$ where for $\sigma \subseteq [n] \times [n]$

$$R_i(\sigma) = \begin{cases} \sqrt{A_{i,j}} & \text{if } |\sigma \cap (\{i\} \times [n])| = 1 \text{ and} \\ & \sigma \cap (\{i\} \times [n]) = \{(i, j)\}, \\ 0 & \text{if } |\sigma \cap (\{i\} \times [n])| \neq 1, \end{cases}$$

$$C_j(\sigma) = \begin{cases} \sqrt{A_{i,j}} & \text{if } |\sigma \cap ([n] \times \{j\})| = 1 \text{ and} \\ & \sigma \cap ([n] \times \{j\}) = \{(i, j)\}, \\ 0 & \text{if } |\sigma \cap ([n] \times \{j\})| \neq 1. \end{cases}$$

With this definition we have

$$\left(\prod_i R_i(\sigma) \right) \left(\prod_j C_j(\sigma) \right) = \begin{cases} \prod_{e \in \sigma} A_e & \text{if } \sigma \text{ is a permutation,} \\ 0 & \text{otherwise.} \end{cases}$$

A practically successful heuristic for computing the partition function of graphical models is Bethe approximation; in fact for "loopless" graphical models the Bethe approximation exactly computes the partition function. The graphical model defined here is not "loopless", but the Bethe approximation can still be defined and in fact computed efficiently. For a detailed account of these, see Vontobel [35]. We only describe the Bethe approximation for the aforementioned graphical model here, but we remark that Bethe approximation can be defined for arbitrary graphical models.

Definition 6. For a matrix $A \in \mathbb{R}_{\geq 0}^{n \times n}$ and a doubly stochastic matrix $P \in \mathcal{B}_n$, define

$$\beta(A, P) = \sum_{e \in [n] \times [n]} (P_e \log(A_e/P_e) + (1 - P_e) \log(1 - P_e)).$$

Note that according to this definition and definition 2 we have

$$\text{bethe}(A) = \max\{e^{\beta(A, P)} \mid P \in \mathcal{B}_n\}.$$

Vontobel [35] proved that $\beta(A, P)$ is a concave function of P . Gurvits [21] proved that $\beta(A, P)$ is always a lower bound for $\log(\text{per}(A))$ by using an inequality of Schrijver [31]. For two recent alternative proofs see Csikvári [16] and Anari and Oveis Gharan [5].

The matrix P maximizing $\beta(A, P)$ is not necessarily obtained from A by matrix scaling. It can be either obtained by convex programming methods since $\beta(A, P)$ is concave in P , or by the sum-product algorithm [35]. Despite this, Gurvits and Samorodnitsky [22] analyzed the objective $\beta(A, P)$ for the matrix P obtained from matrix scaling; in other words they analyzed $\beta(A, RAC)$ where R, C are diagonal matrices such that $RAC \in \mathcal{B}_n$. They showed that

$$\text{per}(A) \leq 2^n e^{\beta(A, RAC)}.$$

We choose a different doubly stochastic matrix. We analyze $\beta(A, P)$ where P is the doubly stochastic matrix of μ 's marginals. In other words we let $P_{i,j} = \mathbb{P}_{\sigma \sim \mu}[(i, j) \in \sigma]$, where $\mu(\sigma) \propto \prod_{e \in \sigma} A_e$.

Theorem 7. For P defined as above, we have

$$\text{per}(A) \leq \sqrt{2}^n e^{\beta(A, P)}.$$

Note that theorem 7 immediately implies theorem 4. We will prove theorem 7 in section IV, but in the rest of this section we demonstrate why the factor $\sqrt{2}^n$ is tight. The example we use was already studied in [22], but for completeness we mention it here as well.

First let us compute $\beta(A, P)$, where A is a 2×2 matrix. Note that \mathcal{B}_2 is a one-dimensional family that can be described by a single parameter $p \in [0, 1]$:

$$\mathcal{B}_2 = \left\{ \begin{bmatrix} p & 1-p \\ 1-p & p \end{bmatrix} \mid p \in [0, 1] \right\}.$$

Then it is an easy exercise to see $\beta(A, P) = p \log(A_{1,1}A_{2,2}) + (1-p) \log(A_{1,2}A_{2,1})$, and as such $\text{bethe}(A) = \max\{A_{1,1}A_{2,2}, A_{1,2}A_{2,1}\}$. Clearly there can be a factor of 2 difference between $\max\{A_{1,1}A_{2,2}, A_{1,2}A_{2,1}\}$ and the sum, $\text{per}(A) = A_{1,1}A_{2,2} + A_{1,2}A_{2,1}$. This happens for example for the matrix

$$A = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix},$$

for which $\text{per}(A) = 2 \text{bethe}(A) = \sqrt{2}^2 \text{bethe}(A)$. This shows that $\sqrt{2}^n$ is tight for $n = 2$.

For higher values of n , note that both $\text{per}(\cdot)$ and $\text{bethe}(\cdot)$ behave predictably w.r.t. block-diagonal matrices. In other words if $B \in \mathbb{R}_{\geq 0}^{n_1 \times n_1}$ and $C \in \mathbb{R}_{\geq 0}^{n_2 \times n_2}$, and A is defined as

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

then $\text{per}(A) = \text{per}(B)\text{per}(C)$ and $\text{bethe}(A) = \text{bethe}(B)\text{bethe}(C)$. To see the latter, note that if $P \in \mathcal{B}_{n_1+n_2}$ has any nonzero entry in places where A has a zero entry, then $\beta(A, P) = -\infty$. So to compute $\text{bethe}(A)$, it is enough to maximize $\beta(A, P)$ over doubly stochastic matrices P which also have a block-diagonal structure

$$P = \begin{bmatrix} Q & 0 \\ 0 & R \end{bmatrix}.$$

But $P \in \mathcal{B}_n$ iff $Q \in \mathcal{B}_{n_1}$ and $R \in \mathcal{B}_{n_2}$. This shows that $\beta(A, P) = \beta(B, Q) + \beta(C, R)$ and consequently $\text{bethe}(A) = \text{bethe}(B)\text{bethe}(C)$. So starting from any matrix, we can construct higher-dimensional matrices by using a block-diagonal form and amplify the gap between $\text{per}(\cdot)$ and $\text{bethe}(\cdot)$. In particular for the $n \times n$ matrix

$$A = I_{n/2} \otimes \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 & 0 & \dots & 0 \\ 1 & 1 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 1 & \dots & 0 \\ 0 & 0 & 1 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1 \end{bmatrix},$$

we have $\text{per}(A) = 2^{n/2} = \sqrt{2}^n$ and $\text{bethe}(A) = 1$. Finally, we remark that it is also easy to construct examples A where $\text{per}(A) = \text{bethe}(A)$. For example this happens when A is a diagonal matrix, such as the identity matrix.

IV. ENTROPY-BASED UPPER BOUND ON THE PERMANENT

In this section we prove theorem 7. Our strategy is to upper bound the permanent by comparing two distributions on perfect matchings in terms of the KL divergence. We will ultimately show that the gap in our upper bound is related to the maximum of a function $\phi : \Delta_n \rightarrow \mathbb{R}$ which is independent of the matrix A . Then, in section V we analyze this function and obtain its maximum. We find it instructive to have the tight example from the end of section III in mind throughout the proof and notice that the inequalities used are tight for this example.

Let $A \in \mathbb{R}_{\geq 0}^{n \times n}$ be a fixed matrix. Since we are upper bounding the permanent by a nonnegative quantity, we can assume w.l.o.g. that $\text{per}(A) > 0$. Then we can define a distribution μ on S_n by $\mu(\sigma) \propto \prod_{e \in \sigma} A_e$, or in other words

$$\mu(\sigma) = \frac{\prod_{e \in \sigma} A_e}{\text{per}(A)}.$$

We will also fix a doubly stochastic matrix P which is defined by

$$P_{i,j} = \mathbb{P}_{\sigma \sim \mu}[(i, j) \in \sigma].$$

P is doubly stochastic because for each $\sigma \in S_n$ and every i there is a unique j such that $(i, j) \in \sigma$, and for every j there is a unique i such that $(i, j) \in \sigma$.

Now let us fix an ordering $\pi \in S_n$ over rows of P and define another distribution ν_π on S_n by the following sampling procedure: We iterate over rows of P according to the order π . In iteration i we select a yet unselected column j with probability proportional to $P_{\pi(i),j}$. We let σ be the set of all $(\pi(i), j)$ selected by this procedure.

Let us compute $\nu_\pi(\sigma)$ explicitly for some given σ . When the sampling procedure produces $\sigma \in S_n$, in iteration i , the

unselected columns are exactly $\sigma \circ \pi(i), \sigma \circ \pi(i+1), \dots, \sigma \circ \pi(n)$. So we pick the first one, $\sigma \circ \pi(i)$, with probability

$$\frac{P_{\pi(i), \sigma \circ \pi(i)}}{\sum_{j=i}^n P_{\pi(i), \sigma \circ \pi(j)}}.$$

This means that

$$\nu_\pi(\sigma) = \prod_{i=1}^n \frac{P_{\pi(i), \sigma \circ \pi(i)}}{\sum_{j=i}^n P_{\pi(i), \sigma \circ \pi(j)}} = \prod_{i=1}^n \frac{P_{i, \sigma(i)}}{\sum_{j \geq \sigma \circ \pi(i)} P_{i, j}}, \quad (1)$$

where as a reminder $j \geq_{\sigma \circ \pi} \sigma(i)$ means that $j = \sigma(i)$ or j appears after $\sigma(i)$ in the list $\sigma \circ \pi(1), \dots, \sigma \circ \pi(n)$.

It is worth mentioning that if A was the tight example presented at the end of [section III](#), regardless of the choice of π , we would have $\mu = \nu_\pi$. But this is generally not true and μ can be quite different from ν_π . Regardless, by [fact 5](#), we have

$$\mathcal{D}(\mu \parallel \nu_\pi) = \mathbb{E}_{\sigma \sim \mu} \left[\log \frac{\mu(\sigma)}{\nu_\pi(\sigma)} \right] \geq 0. \quad (2)$$

Now we decompose [eq. \(2\)](#) part by part. First we have

$$\begin{aligned} \mathbb{E}_{\sigma \sim \mu} [\log \mu(\sigma)] &= \mathbb{E}_{\sigma \sim \mu} \left[\log \frac{\prod_{e \in \sigma} A_e}{\text{per}(A)} \right] \\ &= \sum_{e \in [n] \times [n]} \mathbb{P}_{\sigma \sim \mu} [e \in \sigma] \log(A_e) \\ &\quad - \log(\text{per}(A)) \\ &= \sum_{e \in [n] \times [n]} P_e \log(A_e) - \log(\text{per}(A)), \end{aligned} \quad (3)$$

where the last equality follows from our choice of P , for which we have $\mathbb{P}_{\sigma \sim \mu} [e \in \sigma] = P_e$. For the second part we use [eq. \(1\)](#) to derive

$$\begin{aligned} \mathbb{E}_{\sigma \sim \mu} [\log \nu_\pi(\sigma)] &= \mathbb{E}_{\sigma \sim \mu} \left[\log \left(\prod_{i=1}^n \frac{P_{i, \sigma(i)}}{\sum_{j \geq_{\sigma \circ \pi} \sigma(i)} P_{i, j}} \right) \right] \\ &= \sum_{e \in [n] \times [n]} P_e \log(P_e) \\ &\quad - \sum_{i \in [n]} \mathbb{E}_{\sigma \sim \mu} \left[\log \sum_{j \geq_{\sigma \circ \pi} \sigma(i)} P_{i, j} \right] \end{aligned} \quad (4)$$

Combining [eqs. \(2\)](#) to [\(4\)](#) we get the inequality

$$\sum_{e \in [n] \times [n]} P_e \log \left(\frac{A_e}{P_e} \right) + \sum_{i \in [n]} \mathbb{E}_{\sigma \sim \mu} \left[\log \sum_{j \geq_{\sigma \circ \pi} \sigma(i)} P_{i, j} \right] \geq \log(\text{per}(A)). \quad (5)$$

This already looks very promising as we have $\log(\text{per}(A))$ on one side and part of $\beta(A, P)$ on the other – see [definition 6](#). However the extra sum does not look easy to analyze. In order to simplify it, we take the expectation of [eq. \(5\)](#) over

π chosen *uniformly* at random from S_n , which we denote by $\pi \sim S_n$. Then we get

$$\begin{aligned} &\sum_{e \in [n] \times [n]} P_e \log(A_e/P_e) + \\ &\sum_{i \in [n]} \mathbb{E}_{\sigma \sim \mu} \left[\mathbb{E}_{\pi \sim S_n} \left[\log \sum_{j \geq_{\sigma \circ \pi} \sigma(i)} P_{i, j} \right] \right] \geq \\ &\log(\text{per}(A)). \end{aligned}$$

The main benefit here is that for any σ , when π is a uniformly random element of S_n , $\sigma \circ \pi$ is also a uniformly random element of S_n . So we can simplify this to

$$\begin{aligned} &\sum_{e \in [n] \times [n]} P_e \log(A_e/P_e) + \\ &\sum_{i \in [n]} \mathbb{E}_{\sigma \sim \mu} \left[\mathbb{E}_{\pi \sim S_n} \left[\log \sum_{j \geq_{\pi} \sigma(i)} P_{i, j} \right] \right] \geq \\ &\log(\text{per}(A)). \end{aligned}$$

Notice that we have removed all dependence on σ except for $\sigma(i)$. So we can simplify and get

$$\begin{aligned} &\sum_{e \in [n] \times [n]} P_e \log(A_e/P_e) + \\ &\sum_{i \in [n]} \sum_{k \in [n]} P_{i, k} \cdot \mathbb{E}_{\pi \sim S_n} \left[\log \sum_{j \geq_{\pi} k} P_{i, j} \right] \geq \\ &\log(\text{per}(A)). \end{aligned}$$

Now all appearances of μ have been removed. Our goal is to show that $\beta(A, P) + n \log(\sqrt{2}) \geq \log(\text{per}(A))$. Comparing with the above and [definition 6](#), it is enough to prove the following for each $i = 1, \dots, n$:

$$\begin{aligned} &\log(\sqrt{2}) + \sum_{k \in [n]} (1 - P_{i, k}) \log(1 - P_{i, k}) \geq \\ &\sum_{k \in [n]} P_{i, k} \cdot \mathbb{E}_{\pi \sim S_n} \left[\log \sum_{j \geq_{\pi} k} P_{i, j} \right]. \end{aligned}$$

Proving this would finish the proof of [theorem 7](#). Note that this is a statement that only involves the i -th row of P . So we can abstract away and prove the following lemma.

Lemma 8. *Suppose that $(p_1, \dots, p_n) \in \Delta_n$. Then*

$$\begin{aligned} &\log(\sqrt{2}) \geq \\ &\mathbb{E}_{\pi \sim S_n} \left[\sum_{k=1}^n \left(p_k \log \left(\sum_{j \geq_{\pi} k} p_j \right) - (1 - p_k) \log(1 - p_k) \right) \right]. \end{aligned}$$

Proof: Ideally we would like to prove the inequality for each $\pi \in S_n$ and remove the expectation over π . However, the inequality does not hold individually for each $\pi \in S_n$.

Instead we pair up permutations π with their reverses $\bar{\pi}$, and prove that the sum of the value inside the r.h.s. expectation over π and $\bar{\pi}$ for any fixed π is at most twice the l.h.s. This would finish the proof. Note that the only appearance of π on the r.h.s. is in the ordering \geq_π it induces. So up to reindexing of p_1, \dots, p_n it is enough to prove this for the identity permutation, i.e., it is enough to prove

$$\begin{aligned} \log(2) \geq & \sum_{k \in [n]} p_k \log(p_1 + \dots + p_k) + \\ & \sum_{k \in [n]} p_k \log(p_k + \dots + p_n) - \\ & 2 \sum_{k \in [n]} (1 - p_k) \log(1 - p_k). \end{aligned}$$

Naming the r.h.s. of the above as $\phi(p_1, \dots, p_n)$, we will prove in [section V](#), in [lemma 10](#), that $\max\{\phi(p) \mid p \in \Delta_n\}$ is at most $\log(2)$. This finishes the proof since

$$\begin{aligned} \mathbb{E}_{\pi \sim S_n} \left[\sum_{k=1}^n \left(p_k \log \left(\sum_{j \geq \pi^k} p_j \right) - (1 - p_k) \log(1 - p_k) \right) \right] \\ = \frac{1}{2} \mathbb{E}_{\pi \sim S_n} [\phi(p_{\pi(1)}, \dots, p_{\pi(n)})] \leq \log(\sqrt{2}). \end{aligned}$$

V. ANALYZING THE MAXIMUM GAP

In this section we prove that the maximum of the function $\phi : \Delta_n \rightarrow \mathbb{R}$ is $\log(2)$, finishing the proof of [theorem 7](#). As a reminder, we define the function ϕ again below.

Definition 9. For $p = (p_1, \dots, p_n) \in \Delta_n$ define

$$\begin{aligned} \phi(p_1, \dots, p_n) = & \sum_{k \in [n]} p_k \log(p_1 + \dots + p_k) + \\ & \sum_{k \in [n]} p_k \log(p_k + \dots + p_n) - \\ & 2 \sum_{k \in [n]} (1 - p_k) \log(1 - p_k). \end{aligned}$$

With some abuse of notation we use the same letter ϕ to refer to this function over Δ_n for any n .

A simple calculation shows that

$$\phi(1/2, 1/2) = \log(2).$$

Also notice that zero coordinates in the input to ϕ can be dropped without changing the value of ϕ . In other words, for any $p \in \Delta_{n-1}$ and any i , we have

$$\phi(p_1, \dots, p_i, 0, p_{i+1}, \dots, p_{n-1}) = \phi(p_1, \dots, p_{n-1}).$$

This shows that for any two distinct elements $\{i, j\} \subseteq [n]$ we have

$$\phi(\mathbb{1}_{\{i,j\}}/2) = \log(2).$$

Our goal is to show that these $\binom{n}{2}$ points are the global maxima of ϕ over Δ_n . The main result of this section is

Lemma 10. For any $p \in \Delta_n$ we have $\phi(p) \leq \log(2)$.

Our strategy for proving [lemma 10](#) is to prove it for $n = 1, 2, 3$ and then reduce general n to one of these cases. The main tool enabling this reduction is the following lemma.

Lemma 11. Suppose that $(q, r, s, t) \in \Delta_4$ and $r + s \leq \gamma$ where $\gamma = (\sqrt{17} - 3)/2 \simeq 0.56$ is a constant. Then

$$\phi(q, r, s, t) \leq \phi(q, r + s, t).$$

Let us now see why [lemma 10](#) can be proved, assuming the cases $n = 1, 2, 3$ and [lemma 11](#).

Proof of lemma 10 assuming the proof for $n = 1, 2, 3$ and lemma 11: Suppose that $p = (p_1, \dots, p_n) \in \Delta_n$. If there is an i such that $p_i + p_{i+1} \leq \gamma$, then we claim that

$$\phi(p_1, \dots, p_n) \leq \phi(p_1, \dots, p_{i-1}, p_i + p_{i+1}, p_{i+2}, \dots, p_n).$$

To see this let $q = p_1 + \dots + p_{i-1}$, $r = p_i$, $s = p_{i+1}$, and $t = p_{i+2} + \dots + p_n$. Notice that $(q, r, s, t) \in \Delta_4$ and it follows from [definition 9](#) that

$$\begin{aligned} & \phi(p_1, \dots, p_{i-1}, p_i + p_{i+1}, p_{i+2}, \dots, p_n) - \\ & \phi(p_1, \dots, p_n) = \\ & \phi(q, r + s, t) - \phi(q, r, s, t), \end{aligned}$$

and by [lemma 11](#) this quantity is nonnegative.

We just proved that if there is a consecutive pair p_i, p_{i+1} in p_1, \dots, p_n that add to at most γ we can reduce n without decreasing ϕ , by replacing p_i, p_{i+1} with $p_i + p_{i+1}$. Notice that such a pair always exists for $n \geq 4$. This is because $(p_1 + p_2) + (p_3 + p_4) \leq 1$, so at least one of $p_1 + p_2$ and $p_3 + p_4$ must be $\leq 1/2 < \gamma$. So by induction the problem gets reduced to the cases $n = 1, 2, 3$. ■

Next we prove [lemma 11](#). Afterwards we prove [lemma 10](#) for $n = 1, 2, 3$ finishing the whole proof.

Proof of lemma 11: We want to show that $\phi(q, r, s, t) - \phi(q, r + s, t) \leq 0$. To reduce the number of variables, let us fix r, s and choose q, t in a way that $\phi(q, r, s, t) - \phi(q, r + s, t)$ is maximized. It is then enough to show that the maximum is nonpositive. One can compute that

$$\begin{aligned} & \phi(q, r, s, t) - \phi(q, r + s, t) = \\ & r \log \left(1 - \frac{s}{q + r + s} \right) + s \log \left(1 - \frac{r}{r + s + t} \right) - \\ & 2(1 - r) \log(1 - r) - 2(1 - s) \log(1 - s) + \\ & 2(1 - r - s) \log(1 - r - s). \end{aligned}$$

If we replace q, t by $q + \epsilon, t - \epsilon$ for some $\epsilon \in [-q, t]$, then only the first two terms in the above change. One can easily

check that the above is a concave function of ϵ by checking concavity of the first two summands in terms of q, t . So in order to find the ϵ maximizing the above, we can set the derivative w.r.t. ϵ to zero. If q, t were already the maximizers, then either the derivative at $\epsilon = 0$ would vanish or one of q, t would be 0. We will show that the first case happens for $r + s \leq \gamma$, and obtain an explicit expression for the maximizing q, t . Let us compute the derivative:

$$\frac{d}{d\epsilon}(\phi(q + \epsilon, r, s, t - \epsilon) - \phi(q + \epsilon, r + s, t - \epsilon)) \Big|_{\epsilon=0} = \frac{s}{(q+r)(q+r+s)} - \frac{r}{(s+t)(r+s+t)}.$$

One can compute that this derivative vanishes at

$$\begin{aligned} q^*(r, s) &= \frac{1 - r(1 + r + s)}{2 + r + s} \\ t^*(r, s) &= \frac{1 - s(1 + r + s)}{2 + r + s}, \end{aligned} \quad (6)$$

and for this choice of $q = q^*(r, s), t = t^*(r, s)$ we have $q + r + s + t = 1$. So as long as $q^*(r, s), t^*(r, s) \geq 0$, this will be the maximizer of $\phi(q, r, s, t) - \phi(q, r + s, t)$ for fixed r, s . But $q^*(r, s), t^*(r, s) \geq 0$ exactly when $(1 + r + s) \cdot \max(r, s) \leq 1$. Since we have assumed $s + t \leq \gamma$, we can deduce

$$(1 + r + s) \max(r, s) \leq (1 + \gamma)\gamma < 1,$$

where the last inequality can be easily checked for the constant $\gamma = (\sqrt{17} - 3)/2$. So from now on, we assume that eq. (6) holds and $q = q^*(r, s)$ and $t = t^*(r, s)$. With this choice we can define

$$\begin{aligned} \psi(r, s) &:= \\ \phi(q^*(r, s), r, s, t^*(r, s)) - \phi(q^*(r, s), r + s, t^*(r, s)) &= \\ - (r + s) \log(1 + r + s) + (s - r) \log\left(\frac{1 + s}{1 + r}\right) & \\ - 2(1 - r) \log(1 - r) - 2(1 - s) \log(1 - s) + & \\ 2(1 - r - s) \log(1 - r - s). & \end{aligned}$$

Our goal is to prove that $\psi(r, s) \leq 0$ as long as $r + s \leq \gamma$. Let us call $C = r + s$. We can look at the function $\psi(x, C - x)$ as x varies from 0 to C ; then our goal would be to show $x \mapsto \psi(x, C - x)$ is nonpositive over this range, as long as $C \leq \gamma$. Note that $\psi(\cdot, \cdot)$ is symmetric in its inputs, so it is enough to prove $x \mapsto \psi(x, C - x)$ is nonpositive for $x \in [0, C/2]$. For this we will show that for $x \in [0, C/2]$ we have $\frac{d}{dx} \psi(x, C - x) \leq 0$. This would finish the proof because assuming the derivative is nonpositive, we have $\psi(0, C) = 0$ and

$$\psi(x, C - x) = \int_0^x \frac{d}{dy} \psi(y, C - y) dy \leq 0.$$

So all that remains to finish the proof is to show that $\frac{d}{dx} \psi(x, C - x) \leq 0$ when $x \in [0, C/2]$ and $C \leq \gamma$. For

brevity of expressions we use $r = x$ and $s = C - x$ and notice that $\frac{d}{dx} r = 1$ and $\frac{d}{dx} s = -1$. The term $r + s$ and functions of $r + s$ are constant w.r.t. x , so their derivatives vanish. This simplifies the calculations and we obtain $\frac{d}{dx} \psi(r, s) = \frac{d}{dx} \psi(x, C - x) =$

$$\begin{aligned} \frac{d}{dx} \left((s - r) \log\left(\frac{1 + s}{1 + r}\right) - 2(1 - r) \log(1 - r) - \right. & \\ \left. 2(1 - s) \log(1 - s) \right) = & \\ - 2 \log\left(\frac{1 + s}{1 + r}\right) - (s - r) \left(\frac{1}{1 + r} + \frac{1}{1 + s} \right) + & \\ 2 \log\left(\frac{1 - r}{1 - s}\right) = & \\ (s - r) \cdot \frac{2 \log\left(\frac{1}{1 - z^2}\right) \Big|_{z=r}^{z=s}}{s - r} - & \\ (s - r) \cdot \left(\frac{1}{1 + r} + \frac{1}{1 + s} \right). & \end{aligned}$$

Noting that $s - r \geq 0$, it is enough to prove that

$$\frac{2 \log\left(\frac{1}{1 - z^2}\right) \Big|_{z=r}^{z=s}}{s - r} \leq \frac{1}{1 + r} + \frac{1}{1 + s}.$$

Note that $\frac{d}{dz} (2 \log(1/(1 - z^2))) = 4z/(1 - z^2)$. So we can rewrite the above as

$$\frac{\int_r^s \frac{4z}{1 - z^2} dz}{s - r} \leq \frac{1}{1 + r} + \frac{1}{1 + s}.$$

The l.h.s. is the average value of $4z/(1 - z^2)$ over the interval $[r, s]$. The function $z \mapsto 4z/(1 - z^2)$ is convex over $[r, s] \subseteq [0, 1]$, so the average over $[r, s]$ can be upper bounded by the average over the two endpoints $\{r, s\}$. So it would be enough to prove

$$\frac{2r}{1 - r^2} + \frac{2s}{1 - s^2} \leq \frac{1}{1 + r} + \frac{1}{1 + s}.$$

Multiplying all sides by $(1 - r^2)(1 - s^2)$ and rearranging the terms it would be enough to prove the following

$$(1 - 3r)(1 - s^2) + (1 - 3s)(1 - r^2) \geq 0.$$

The l.h.s. can be rewritten as

$$\begin{aligned} 2 - 3(r + s)(1 - rs) - (r^2 + s^2) &\geq 2 - 3(r + s) - (r + s)^2 = \\ 2 - 3C - C^2, & \end{aligned}$$

where for the inequality we used the fact that $1 - rs \leq 1$ and $r^2 + s^2 \leq (r + s)^2$. The above is a quadratic in C which has roots at $\frac{1}{2}(\pm\sqrt{17} - 3)$. So it is nonnegative between the roots, and in particular on the interval $C \in [0, (\sqrt{17} - 3)/2] = [0, \gamma]$. ■

Now we just need to prove lemma 10 for $n = 1, 2, 3$.

Proof of lemma 10 for $n = 1$: For $n = 1$, the only $p \in \Delta_1$ is $p = 1$. We can easily compute $\phi(1) = 0 \leq \log(2)$ as claimed. ■

Proof of lemma 10 for $n = 2$: We have $\Delta_2 = \{(q, 1 - q) \mid q \in [0, 1]\}$. So we can calculate and see that

$$\phi(q, 1 - q) = -q \log(q) - (1 - q) \log(1 - q).$$

This is a concave function of q and it is symmetric on $[0, 1]$. So its maximum is achieved at $q = 1/2$, where it takes the value $\phi(1/2, 1/2) = \log(2)$. ■

Proof of lemma 10 for $n = 3$: We have to show that if $q, r, s \geq 0$ and $q + r + s = 1$, then $\phi(q, r, s) \leq \log(2)$. This looks difficult to prove at first; there are at least three points where equality can be attained, namely $(1/2, 0, 1/2)$, $(1/2, 1/2, 0)$, $(0, 1/2, 1/2)$, and we want to prove these are the maxima. This suggests that convexity arguments would not work, and we need to somehow break the symmetry. One can gain some confidence about the correctness of the statement by using a computer program to calculate ϕ over an ϵ -net of Δ_3 , but this would not constitute a proof. Instead we use some tricks to reduce the search space to a finite set and use exact rational arithmetic, in order to formally verify the statement with the help of a computer program.

Our insight is that we only need to upper bound $\phi(q, r, s)$ for $(q, r, s) \in U \subseteq \Delta_3$, where U is a subset of Δ_3 which is far from the three points conjectured to be the maxima. To see this, note that if $q + r \leq \gamma$ or $r + s \leq \gamma$, then we can again use lemma 11 to reduce to the case $n = 2$ which we have already proved. So we can define

$$U = \{(q, r, s) \in \Delta_3 \mid q + r \geq \gamma \wedge r + s \geq \gamma\} = \{(q, r, s) \in \Delta_3 \mid q, s \leq 1 - \gamma\}.$$

Notice that none of the three conjectured maxima are inside U as for each one of them either $q = 1/2 > 1 - \gamma$ or $s = 1/2 > 1 - \gamma$.

So now we can potentially form an ϵ -net for U , bound the error of approximating ϕ at any point by the closest net point and hope that the approximation error is smaller than $\log(2) - \max\{\phi(p) \mid p \in U\}$. The latter would be a positive number, so taking ϵ small enough should in principle yield a proof.

The main problem with this approach would be computing ϕ over the ϵ -net since we would potentially be dealing with irrational numbers that we cannot even write down. We can however resolve this by taking exp from both sides of the desired inequality $\phi(q, r, s) \leq \log(2)$. Let us first simplify the expression for $\phi(q, r, s)$ by replacing r with $1 - q - s$ and calculating

$$\phi(q, r, s) = q \log(q) + s \log(s) - (1 - q + s) \log(1 - q) - (1 + q - s) \log(1 - s) - 2(q + s) \log(q + s).$$

Taking exp we get

$$\exp(\phi(q, r, s)) = \frac{q^q s^s}{(1 - q)^{1 - q + s} (1 - s)^{1 + q - s} (q + s)^{2q + 2s}}.$$

So if we want to show that $\phi(q, r, s) \leq \log(2)$, it is enough to show that the above is at most 2. However note that even if q, s were rational numbers, the above would not necessarily be rational. We again use a simple trick of raising to a power. Suppose that $N \in \mathbb{Z}_{>0}$ is a natural number. Then $\phi(q, r, s) \leq \log(2)$ iff

$$q^{Nq} s^{Ns} \leq 2^N (1 - q)^{N(1 - q + s)} (1 - s)^{N(1 + q - s)} (q + s)^{2N(q + s)}.$$

Now if q, s are rationals and N is such that $Nq, Ns \in \mathbb{Z}$, then both sides of the above would be rational numbers. So we can potentially compute both sides exactly and compare them.

Combining all the ideas, suppose that N is a natural number and let $\epsilon = 1/N$. Suppose that $\bar{q} = i/N$ and $\bar{s} = j/N$ for some $i, j \in \{0, \dots, N - 1\}$ and we want to certify that $\phi(q, 1 - q - s, s) \leq \log(2)$ for all $(q, s) \in [\bar{q}, \bar{q} + \epsilon] \times [\bar{s}, \bar{s} + \epsilon]$. Then it would suffice to check the following

$$(\bar{q} + \epsilon)^i (\bar{s} + \epsilon)^j \leq 2^N (1 - \bar{q} - \epsilon)^{N - i + j + 1} (1 - \bar{s} - \epsilon)^{N + i - j + 1} (\bar{q} + \bar{s})^{2(i + j + 2)}, \quad (7)$$

where we have upper bounded the terms on the l.h.s. and lower bounded the terms on the r.h.s. All numbers appearing as the base of the exponent are at most 1, so on the l.h.s. we want the minimum possible exponents and on the r.h.s. the opposite.

Our goal is to cover U with patches of the form $[\bar{q}, \bar{q} + \epsilon] \times [\bar{s}, \bar{s} + \epsilon]$ and for each such \bar{q}, \bar{s} verify eq. (7). One might hope that taking N large enough would give us enough room to absorb the approximation errors. Unfortunately there is a corner case that never gets resolved for any N ; eq. (7) will never be satisfied for $\bar{q}, \bar{s} = 0$, as the r.h.s. will be 0. The problematic term is the last one: $(\bar{q} + \bar{s})^{2(i + j + 2)}$. To resolve this, note that x^x is a monotone decreasing function of x as long as $x < 1/e$. So if $\bar{q} + \bar{s} + 2\epsilon < 1/e$, then we can replace the problematic term with $(\bar{q} + \bar{s} + 2\epsilon)^{2(i + j + 2)}$ which is nonzero. The condition $\bar{q} + \bar{s} + 2\epsilon < 1/e$ is automatically satisfied for $\bar{q} = \bar{s} = 0$ when $N > 2e$.

The algorithm described above can be seen in algorithm 1. The arithmetic has to be done exactly by using fractional numbers. We implemented the algorithm in Python 3, which supports infinite precision arithmetic with fractional numbers natively, and found that for $N = 2000$ all tests pass successfully. The implementation can be found in algorithm 2, which finished running in a few hours on a MacBook Pro computer. ■

REFERENCES

- [1] Scott Aaronson and Alex Arkhipov. “The computational complexity of linear optics”. In: *Proceedings of the forty-third annual ACM symposium on Theory of computing*. ACM, 2011, pp. 333–342.

Algorithm 1: An algorithm to verify $\max\{\phi(q, r, s) \mid (q, r, s) \in U\} \leq \log(2)$.

```

 $\epsilon \leftarrow 1/N$ 
 $M \leftarrow \lfloor 44N/100 \rfloor$  /* This is the upper bound for  $i, j$ , because  $1 - \gamma < 0.44$ . */
for  $i = 0, 1, \dots, M$  do
  for  $j = 0, 1, \dots, M$  do
     $\bar{q} \leftarrow \frac{i}{N}$ 
     $\bar{s} \leftarrow \frac{j}{N}$ 
    if  $i = j = 0$  and  $N > 100$  then
      Verify the following:
      
$$(\bar{q} + \epsilon)^i (\bar{s} + \epsilon)^j \leq 2^N (1 - \bar{q} - \epsilon)^{N-i+j+1} (1 - \bar{s} - \epsilon)^{N+i-j+1} (\bar{q} + \bar{s} + 2\epsilon)^{2(i+j+2)}$$

    else
      Verify the following:
      
$$(\bar{q} + \epsilon)^i (\bar{s} + \epsilon)^j \leq 2^N (1 - \bar{q} - \epsilon)^{N-i+j+1} (1 - \bar{s} - \epsilon)^{N+i-j+1} (\bar{q} + \bar{s})^{2(i+j+2)}$$

    end
  end
end

```

Algorithm 2: Implementation of algorithm 1 for $N = 2000$ in Python 3.

```

from fractions import Fraction

N = 2000
M = 880 # .44 * N
eps = Fraction(1, N)

for i in range(M+1):
  for j in range(M+1):
    q = Fraction(i, N)
    s = Fraction(j, N)
    LHS = (q+eps)**i * (s+eps)**j
    RHS = 2**N * (1-q-eps)**(N-i+j+1) * (1-s-eps)**(N+i-j+1) * \
          (q+s+(2*eps if i==0 and j==0 else 0))**(2*(i+j+2))
    assert LHS<=RHS

```

- [2] Scott Aaronson and Travis Hance. “Generalizing and derandomizing Gurvits’s approximation algorithm for the permanent”. In: *Quantum Information & Computation* 14.7&8 (2014), pp. 541–559.
- [3] Jayadev Acharya, Hirakendu Das, Alon Orlitsky, and Ananda Theertha Suresh. “A unified maximum likelihood approach for estimating symmetric properties of discrete distributions”. In: *International Conference on Machine Learning*. 2017, pp. 11–21.
- [4] Nima Anari, Leonid Gurvits, Shayan Oveis Gharan, and Amin Saberi. “Simply Exponential Approximation of the Permanent of Positive Semidefinite Matrices”. In: *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, Oct. 2017, pp. 914–925. DOI: 10.1109/focs.2017.89.
- [5] 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*. ACM, June 2017, pp. 384–396. DOI: 10.1145/3055399.3055469.
- [6] Nima Anari, Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. “Nash Social Welfare, Matrix Permanent, and Stable Polynomials”. In: *Proceedings of the 8th Conference on Innovations in Theoretical Computer Science*. Elevated to invited paper. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik,

- Jan. 2017, 36:1–36:12. DOI: 10.4230/LIPIcs.ITCS.2017.36.
- [7] Nima Anari, Shayan Oveis Gharan, and Cynthia Vinzant. “Log-Concave Polynomials, Entropy, and a Deterministic Approximation Algorithm for Counting Bases of Matroids”. In: *Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, Oct. 2018.
- [8] Alexander Barvinok. “Computing permanents of complex diagonally dominant matrices and tensors”. In: *arXiv preprint arXiv:1801.04191* (2018).
- [9] Alexander Barvinok. “Computing the permanent of (some) complex matrices”. In: *Foundations of Computational Mathematics* 16.2 (2016), pp. 329–342.
- [10] Lev M Bregman. “Some properties of nonnegative matrices and their permanents”. In: *Soviet Math. Dokl.* Vol. 14. 4. 1973, pp. 945–949.
- [11] Andrei Z Broder. “How hard is it to marry at random?(On the approximation of the permanent)”. In: *Proceedings of the eighteenth annual ACM symposium on Theory of computing*. ACM. 1986, pp. 50–58.
- [12] L Chakhmakhchyan, NJ Cerf, and R Garcia-Patron. “Quantum-inspired algorithm for estimating the permanent of positive semidefinite matrices”. In: *Physical Review A* 96.2 (2017), p. 022329.
- [13] Moses Charikar, Kirankumar Shiragur, and Aaron Sidford. “Efficient Profile Maximum Likelihood for Universal Symmetric Property Estimation”.
- [14] Sourav Chatterjee and Persi Diaconis. “The sample size required in importance sampling”. In: *The Annals of Applied Probability* 28.2 (2018), pp. 1099–1135.
- [15] Michael Chertkov, Lukas Kroc, and Massimo Vergassola. “Belief propagation and beyond for particle tracking”. In: *arXiv preprint arXiv:0806.1199* (2008).
- [16] Péter Csikvári. “Lower matching conjecture, and a new proof of Schrijver’s and Gurvits’s theorems”. In: *arXiv preprint arXiv:1406.0766* (2014).
- [17] Persi Diaconis. “Sequential importance sampling for estimating the number of perfect matchings in bipartite graphs: An ongoing conversation with Laci”. In: ().
- [18] Lior Eldar and Saeed Mehraban. “Approximating the permanent of a random matrix with vanishing mean”. In: *arXiv preprint arXiv:1711.09457* (2017).
- [19] Daniel Grier and Luke Schaeffer. “New hardness results for the permanent using linear optics”. In: *arXiv preprint arXiv:1610.04670* (2016).
- [20] Leonid Gurvits. “On the complexity of mixed discriminants and related problems”. In: *International Symposium on Mathematical Foundations of Computer Science*. Springer. 2005, pp. 447–458.
- [21] Leonid Gurvits. “Unleashing the power of Schrijver’s permanent inequality with the help of the Bethe approximation”. In: *arXiv preprint arXiv:1106.2844* (2011).
- [22] Leonid Gurvits and Alex Samorodnitsky. “Bounds on the Permanent and Some Applications”. In: *Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society. 2014, pp. 90–99.
- [23] Bert Huang and Tony Jebara. “Approximating the permanent with belief propagation”. In: *arXiv preprint arXiv:0908.1769* (2009).
- [24] Mark Jerrum, Alistair Sinclair, and Eric Vigoda. “A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries”. In: *Journal of the ACM (JACM)* 51.4 (2004), pp. 671–697.
- [25] Nathan Linial, Alex Samorodnitsky, and Avi Wigderson. “A deterministic strongly polynomial algorithm for matrix scaling and approximate permanents”. In: *Combinatorica* 20.4 (2000), pp. 545–568.
- [26] David JC MacKay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.
- [27] Henryk Minc. “Upper bounds for permanents of $(0, 1)$ -matrices”. In: *Bulletin of the American Mathematical Society* 69.6 (1963), pp. 789–791.
- [28] Dmitri S Pavlichin, Jiantao Jiao, and Tsachy Weissman. “Approximate profile maximum likelihood”. In: *arXiv preprint arXiv:1712.07177* (2017).
- [29] Jaikumar Radhakrishnan. “An entropy proof of Bregman’s theorem”. In: *Journal of combinatorial theory, Series A* 77.1 (1997), pp. 161–164.
- [30] Alexander Schrijver. “A short proof of Minc’s conjecture”. In: *Journal of combinatorial theory, Series A* 25.1 (1978), pp. 80–83.
- [31] Alexander Schrijver. “Counting 1-factors in regular bipartite graphs”. In: *J. Comb. Theory, Ser. B* 72.1 (1998), pp. 122–135.
- [32] Damian Straszak and Nisheeth K Vishnoi. “Belief propagation, bethe approximation and polynomials”. In: *Communication, Control, and Computing (Allerton), 2017 55th Annual Allerton Conference on*. IEEE. 2017, pp. 666–671.
- [33] Leslie G Valiant. “The complexity of computing the permanent”. In: *Theoretical computer science* 8.2 (1979), pp. 189–201.
- [34] Pascal O Vontobel. “The Bethe and Sinkhorn approximations of the pattern maximum likelihood estimate and their connections to the Valiant-Valiant estimate.” In: *ITA*. 2014, pp. 1–10.
- [35] Pascal O Vontobel. “The Bethe permanent of a non-negative matrix”. In: *IEEE Transactions on Information Theory* 59.3 (2013), pp. 1866–1901.
- [36] Jonathan S Yedidia, William T Freeman, and Yair Weiss. “Constructing free-energy approximations and

generalized belief propagation algorithms”. In: *IEEE Transactions on information theory* 51.7 (2005), pp. 2282–2312.

- [37] Günter M Ziegler. *Lectures on polytopes*. Vol. 152. Springer Science & Business Media, 2012.