

Dual Charging for Half-Integral TSP

Nathan Klein 

Department of Computer Science, Boston University, MA, USA

Mehrshad Taziki 

Department of Computer Science, ETH Zürich, Zürich, Switzerland

Abstract

In this extended abstract, we show that the max entropy algorithm is a randomized 1.49776 approximation for half-integral TSP, improving upon the previous known bound of 1.49993 from Karlin et al. This also improves upon the best-known approximation for half-integral TSP due to Gupta et al. Our improvement results from using the dual, instead of the primal, to analyze the expected cost of the matching. We believe this method of analysis could lead to a simpler proof that max entropy is a better-than-3/2 approximation in the general case.

2012 ACM Subject Classification Theory of computation → Routing and network design problems

Keywords and phrases Approximation Algorithms, Graph Algorithms, Randomized Rounding, Linear Programming

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2025.21

Category APPROX

1 Introduction

In the metric traveling salesperson problem (TSP), we are given a weighted graph $G = (V, E)$ and aim to find the shortest closed walk that visits every vertex. Metric TSP is NP-Hard to approximate better than $\frac{123}{122}$ [25]. In the 1970s, Christofides and Serdyukov [5, 33] famously gave a $\frac{3}{2}$ approximation for the problem. This was not improved until recently, when Karlin, Klein and Oveis Gharan showed a $\frac{3}{2} - \epsilon$ approximation for $\epsilon = 10^{-36}$ [23, 24] in 2021. Gurvits, Klein, and Leake later showed that one can set $\epsilon = 10^{-34}$ [16].

These recent improvements originate in work from 2010 and 2011 by Asadpour et al. on asymmetric TSP [1] and Oveis Gharan, Saberi and Singh on graphic TSP [29]. These two papers introduced the so-called *max entropy* algorithm for TSP. In this algorithm, one first solves the subtour elimination [6] (or Held-Karp [19]) linear program for TSP to obtain a fractional point $x \in \mathbb{R}_{\geq 0}^E$. Then, the max entropy distribution μ over spanning trees is computed whose marginals match x (up to error ϵ , which can be made exponentially small in n). Finally, a tree T is sampled from μ and the minimum cost matching is added to the odd vertices of T . We know this algorithm is at worst a $\frac{3}{2} - 10^{-34}$ approximation in general, and we also know there are instances where it performs no better than 1.375 [21]. A fascinating open question is to determine what the worst case approximation ratio of the algorithm is.

Prior to the first improvement in the general case, Karlin, Klein and Oveis Gharan showed a randomized 1.49993 approximation for a special case of TSP. In particular, they showed that if for an instance there exists an optimal solution x to the subtour LP with $x_e \in \{0, \frac{1}{2}, 1\}$ for all $e \in E$, then the max entropy algorithm outputs a solution of expected cost at most $1.49993 \cdot c(x)$ [22]. These “half-integral” points are of special interest due to a conjecture of Schalekamp, Williamson, and van Zuylen [31] that the integrality gap of the subtour LP is obtained by half integral points. Notably, the lower bound of 1.375 for max entropy in [21] is a half integral instance, as is the classical envelope graph which demonstrates an integrality gap of at least $\frac{4}{3}$ for the subtour LP (see e.g. [35]).



© Nathan Klein and Mehrshad Taziki;

licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2025).

Editors: Alina Ene and Eshan Chattopadhyay; Article No. 21; pp. 21:1–21:22



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In 2021, Gupta, Lee, Li, Mucha, Newman, and Sarkar claimed to improve the bound for half integral points to 1.4983 [12] using a mix of the max entropy algorithm and a combinatorial one proposed by Haddadan and Newman [18]. After finding a small issue in the proof in [12], we had a discussion with the authors, and while the details are being checked, a fix appears to give a bound of 1.4990.

In this work, we show that the max entropy algorithm (with no adaptations) is a 1.49776 approximation for half-integral TSP. This improves over the state of the art for half integral TSP with a significantly simpler algorithm, as well as gives a large relative improvement for the analysis of max entropy. In particular, we show:

► **Theorem 1.** *Given an optimal solution x to the subtour LP with $x_e \in \{0, \frac{1}{2}, 1\}$ for all $e \in E$, the max entropy algorithm produces a solution of cost at most $1.49776 \cdot c(x)$ in expectation.*

Therefore our result also bounds the integrality gap of the subtour LP in the half integral case by 1.49776 (and, should the conjecture of [31] hold, the general case as well). As discussed in more detail in Section 1.2, the primary reason for the improvement is a new dual-based analysis style. Due to the large relative improvement over the previous analysis of max entropy in the half integral case (in terms of the distance from 1.5, this is an improvement of a factor of about 30) and the fact we no longer need to analyze certain pathological cases, we believe these techniques may lead to a significant simplification of the analysis in the general case.

1.1 Other Related Work

There has been exciting recent progress on two important variants of TSP, *graphic* TSP, in which the input graph is unweighted, and *path* TSP in which the goal is to find the shortest s - t walk visiting every vertex. For graphic TSP, Oveis Gharan, Saberi, and Singh [29] first demonstrated that max entropy was a $\frac{3}{2} - \epsilon$ approximation for a small constant $\epsilon > 0$. Using different methods, Mömke and Svensson [27] then obtained a 1.461 approximation. This was further improved by Mucha [28] to $\frac{13}{9}$ and then to 1.4 by Sebö and Vygen [32]. For path TSP, Zenklusen showed a $\frac{3}{2}$ approximation [36] using a dynamic programming approach. Traub, Vygen, and Zenklusen then showed that given an α approximation for TSP there is an $\alpha + \epsilon$ approximation for Path TSP for any $\epsilon > 0$.

1.2 Overview

As discussed, in the max entropy algorithm we first solve the subtour LP to obtain a solution $x \in \mathbb{R}_{\geq 0}^E$. The subtour LP is as follows, where $\delta(S)$ for $S \subsetneq V$ is the set of edges with exactly one endpoint in S and for $F \subseteq E$, $x(F) = \sum_{e \in F} x_e$.

$$\begin{aligned}
 \min \quad & \sum_{e \in E} c_e x_e \\
 \text{s.t.} \quad & x(\delta(S)) \geq 2 \quad \forall S \subsetneq V \\
 & x(\delta(\{v\})) = 2 \quad \forall v \in V \\
 & x_e \geq 0 \quad \forall e \in E
 \end{aligned} \tag{1}$$

Then, using x , we find a maximum entropy distribution μ over spanning trees¹ subject to the constraint that $\mathbb{P}_{T \sim \mu}[e \in T] = x_e$ for all $e \in E$, up to a small multiplicative error (see Section 2.3 for more discussion on this sampling procedure). Finally, we sample a tree T from μ and add a minimum cost matching M on the odd vertices of the tree T . $T \uplus M$ is an Eulerian graph and thus contains an Eulerian walk, which is a solution to the metric TSP problem.² The main goal, then, is to *bound the expected cost of the matching* over the randomness of the sampled tree T . To do so, we find a function that given a tree T constructs a vector $y \in \mathbb{R}_{\geq 0}^E$ so that $c(M) \leq c(y)$, and then bound the expected cost of y . In particular, y will be a feasible point in the O -Join polytope (where O is the set of odd vertices in T); see Section 2 for further details.

This is the approach taken by all previous papers analyzing max entropy [29, 22, 23] or variants of this algorithm [12], and we do not deviate from this here. However, here we construct y in a new way that streamlines the analysis and allows for a sharper guarantee. While previous works bounded $\mathbb{E}[c(y)]$ by bounding the contribution to y of *each edge individually*, we instead bound the contribution of each dual variable to $\mathbb{E}[c(y)]$. By complementary slackness, the dual variables correspond to tight cuts with $x(\delta(S)) = 2$, which contain multiple edges e with $x_e \in \{\frac{1}{2}, 1\}$. By bounding the cost of *groups* of edges in terms of $\sum_{e \in \delta(S)} y_e$, the argument becomes much more flexible. Previously, analyses had to deal with pathological cases where single edges had a high $\mathbb{E}[y_e]$ and create workarounds. By looking at groups of edges, issues due to these pathological edges are averaged out. Previously it was necessary to bound $\mathbb{E}[y_e] < \frac{x_e}{2}$ for all $e \in E$ (as this would correspond to an expected cost of less than $\frac{1}{2}c(x)$ for the matching), in our construction there may be edges with $\mathbb{E}[y_e] > \frac{x_e}{2}$.

There is one other big advantage of this approach, which is as follows.³ In the analysis of max entropy, typically one begins with $y = \frac{x_e}{2}$ for all edges. Then, depending on the parity of certain cuts in the tree, some edges have y decreased and other edges have y increased. In the per-edge argument, it was important to bound the expected increase of each edge carefully. However, when considering a cut, these increases and decreases often cancel out, simplifying things quite a bit. (Despite this, our analysis is not simpler than [22]. This is because we take care to improve the analysis in several places which require additional casework.)

One other aspect of our improvement is to incrementally improve some of the important probabilistic bounds from [22] using polynomial capacity (see e.g. [16] for usage of this tool in TSP) or more precise arguments. However the impact of this is relatively minor compared to that of the move to a dual-based analysis.

2 Preliminaries

2.1 Notation

Sets and Cuts. Given a set S , let $\delta(S)$ be the set of edges with exactly one endpoint in S and $E(S)$ be the set of edges with both endpoints in S . Given $x \in \mathbb{R}^E$ and $F \subseteq E$, let $x(F) = \sum_{e \in F} x_e$. A set $S \subseteq V$ is *tight* if $x(\delta(S)) = 2$. A set S is *proper* if $2 \leq |S| \leq n - 2$. Two sets A, B *cross* if $A \cap B, A \setminus B, B \setminus A, \overline{A \cup B}$ are all non-empty.

¹ Technically, over spanning trees plus an edge.

² Using the fact that the costs form a metric, one can then shortcut the Eulerian tour to a Hamiltonian cycle of no greater cost if desired.

³ Understanding this advantage is a bit technical, so we recommend readers unfamiliar with work on the max entropy algorithm come back to this point after having read the body of the paper.

21:4 Dual Charging for Half-Integral TSP

Support Graph. We will use x to construct a 4-regular and 4-edge-connected graph which we will call G . G will have the same vertex set as our input. Then for every edge e with $x_e = \frac{1}{2}$, we add the corresponding edge to the support graph. For every edge with $x_e = 1$, we will add two parallel copies of e to G .

Minimum Cuts. Tight sets are therefore minimum cuts of the support graph G , as they have 4 edges. It is helpful to note that since G is Eulerian, every cut which is not a minimum cut has at least 6 edges. For an overview of the structure of minimum cuts, we recommend an unfamiliar reader to look at the cactus representation of Dinits, Karzanov, and Lomonozov [7], or a succinct explanation of it by Frank and Fleiner [10]. This structure is very important to our analysis.

Trees. Given a tree T and a set of edges F , we let $F_T = |T \cap F|$ denote the number of edges of F contained in T . We will use $\delta_T(S)$ to denote the number of edges in $\delta(S)$ contained in T .

Parity. For a sampled tree T , we say a set $S \subseteq V$ is *even* if $\delta_T(S)$ is even and *odd* otherwise.

2.2 Polyhedral Background

The subtour LP is given in Equation (1). We will also crucially use the dual linear program in our analysis. By the parsimonious property [11], the subtour bound does not change after dropping the equality constraints. Thus, the dual can be seen to have the following formulation:

$$\begin{aligned} \max \quad & 2 \sum_{S \subsetneq V} z_S \\ \text{s.t.} \quad & \sum_{S \subsetneq V | e \in \delta(S)} z_S \leq c_e \quad \forall e \in E \\ & z_S \geq 0 \quad \forall S \subsetneq V \end{aligned} \tag{2}$$

We will also make significant use of the following characterization of the cost of the matching. It is well known (see [8], for example) that given a metric, the following LP bounds upper bounds the cost of an integral perfect matching on a set of vertices O :

$$\begin{aligned} \min \quad & \sum_{e \in E} c_e y_e \\ \text{s.t.} \quad & y(\delta(S)) \geq 1 \quad \forall S \subsetneq V, |S \cap O| \text{ odd} \\ & y_e \geq 0 \quad \forall e \in E \end{aligned} \tag{3}$$

This is known as the O -Join polyhedron, and we call a feasible point y in this polyhedron an **O -Join solution**, where O is the set of vertices in the sampled tree T with odd degree. In this context, note that the set of cuts S which have constraints in (3) are exactly those for which $\delta_T(S)$ is odd.

2.3 Max Entropy Distribution

A distribution μ over spanning trees is λ -uniform if $\lambda \in \mathbb{R}_{\geq 0}^E$ and for every spanning tree T of the graph,

$$\mathbb{P}[T] \propto \prod_{e \in T} \lambda_e$$

Given λ , we will let μ_λ denote the resulting λ -uniform distribution. Given a point z in the spanning tree polytope, [1] show that one can efficiently find a λ -uniform tree with marginals arbitrarily close to z :

► **Theorem 2** ([1]). *Let z be a point in the spanning tree polytope of a graph $G = (V, E)$. Then, for any $\epsilon > 0$, there is an algorithm running in time polynomial in the size of the graph and $\log(\frac{1}{\epsilon})$ that outputs a vector $\lambda \in \mathbb{R}_{\geq 0}^E$ so that the λ -uniform distribution μ_λ has the property:*

$$\mathbb{P}_{T \sim \mu_\lambda} [e \in T] \leq (1 + \epsilon) z_e$$

Since ϵ can be chosen to be exponentially small in n , following previous work on half integral TSP, we will assume for brevity that we may set $\epsilon = 0$. This error can be handled using the stability of max entropy distributions [34] (one can see this applied in [23]).

λ -uniform spanning tree distributions have maximum entropy over all distributions with the same marginal vector. Therefore, when one can set $\epsilon = 0$, they are indeed the distributions of maximal entropy respecting the constraints.

2.4 Algorithm and Critical Sets

We will study the version of the max entropy algorithm used by [22]. Here, we first fix an edge $e^+ = (u, v)$ where u, v have two parallel edges between them. Such an edge always exists in any extreme point solution (see [3]), or as noted in [22] one can create such an edge by splitting a vertex in two and putting an edge of value 1 and cost 0 between its two endpoints. After fixing e^+ , we iteratively select a minimal proper tight set S that is not crossed (and $e^+ \notin E(S)$), compute the max entropy distribution μ inside S , sample $T_S \sim \mu$, and contract it. When no such set exists, the graph must be a cycle $v_1, \dots, v_k$, at which point we sample a random cycle. In particular, for every v_i, v_{i+1} that share two edges we sample one of them independently and uniformly at random. The sampled tree T at the end is the union of all trees T_S sampled during the procedure. See Algorithm 1 for a complete description and [22] for further details.

■ **Algorithm 1** Algorithm for half-integral TSP [22].

-
- 1: Let x be a half-integral solution to (1) with an edge e^+ with $x_{e^+} = 1$
 - 2: Set $T = \emptyset$
 - 3: **while** there exists a proper tight set of G that is not crossed (by a tight set) **do**
 - 4: Let S be a minimal such set such that $e^+ \notin E(S)$
 - 5: Compute the maximum entropy distribution μ of $E(S)$
 - 6: Sample a tree from μ and add its edges to T
 - 7: Set $G = G/S$
 - 8: **end while**
 - 9: Sample a uniformly random cycle from G (including e^+) and add it to T
 - 10: Compute the minimum O -Join M on the odd nodes of T and return $T \uplus M$ or its shortcut
-

Following [22], we call every tight set contracted by the algorithm a **critical set**. In addition, we call vertices critical sets. For a critical set S , we call $\delta(S)$ a critical cut. The collection of critical sets is a laminar family $\mathcal{L}$, where recall a family is laminar if for all $S, T \in \mathcal{L}$, $S \cap T \in \{\emptyset, S, T\}$.

There are two types of sets $S \in \mathcal{L}$. If at the moment before contraction, there are no proper minimum cuts inside S , then call S a **degree cut**. Otherwise, call S a **cycle cut**, in which case at the moment before contraction S is a path with two edges between every pair of adjacent vertices. See Figure 1b for examples of each type of cut in G , and see Section 3.1 for more details.

► **Definition 3** (G_S). *For every critical cut S , we let G_S denote the support graph G after contracting every critical cut lower than S in the hierarchy as well as $V \setminus S$ into vertices. We use $G[S]$ to denote the graph $G_S \setminus w$, where w is the vertex representing $V \setminus S$ after contraction.*

We also note here that by definition, edges which are in G_S and $G_{S'}$ for critical cuts $S \neq S'$ are **independent**. This fact will be used crucially throughout the proof.

► **Remark.** We defer additional preliminaries concerning strongly Rayleigh distributions and polynomial capacity to after the proof overview.

3 Overview of our Method

In this section, we introduce the key tools used in our analysis and provide a high-level overview of our techniques. In Section 3.1, we define the necessary notation and describe the structural properties of the hierarchy of critical cuts. Understanding this hierarchy is crucial to our analysis. In Section 3.2, we describe the construction of an O -Join solution with a small expected cost. Furthermore, we provide an explanation of our proof and describe our use of the dual formulation (Equation (2)) in analyzing the cost of the O -Join solution.

3.1 Hierarchy of Critical Cuts

The algorithm constructs a natural hierarchy of min-cuts. That is, the critical cuts form a laminar family of min cuts, and therefore can be arranged in a hierarchical structure. At the bottom of this hierarchy are the singleton vertex cuts.

Let S be a critical cut. Based on the structure of the contracted graph G_S , we classify S into one of two types. Note that, we assumed G_S is the support graph after contracting $V \setminus S$ and all critical cuts lower than S in the hierarchy.

1. **Degree cut:** If G_S contains no proper min-cuts, we call S a degree cut.
2. **Cycle cut:** Otherwise, G_S must form a cycle with two parallel edges between each pair of consecutive vertices. In this case, we call S a cycle cut. We call the parallel edges in $G[S]$ that share their endpoints **companions**. Note that a pair of companions e, f has the property that exactly one is in the tree and the event of which edge is chosen is independent of all other events. The remaining two pairs of parallel edges sharing endpoints in G_S are called **cycle partners**.

► **Fact 4.** *Every min-cut in G is either a critical cut or an interval of a cycle cut.*

For an edge e , let S_e denote the minimal critical cut such that $e \in G[S_e]$. We will distinguish edges into two types depending the structure of S_e .

► **Definition 5** (Top Edge and Bottom Edge). *We call an edge e a **top edge** if S_e is a degree cut, and a **bottom edge** if S_e is a cycle cut.*


We define a similar notation for the critical cuts.

► **Definition 6** (Top Cut and Bottom Cut). A min cut S is said to be a **top cut** if its parent in the hierarchy is a degree cut. Otherwise, S is called a **bottom cut**, i.e., if its parent in the hierarchy is a cycle cut. See Section 3.1 for an example.


Moreover, for an edge e , we define the **Last Cuts** of e as the two maximal min cuts S such that $e \in \delta(S)$. More precisely, let e be a bottom edge where S_e has child cuts $S_1, \dots, S_k$ with two edges between S_i and S_{i+1} for $1 \leq i \leq k-1$. If e is between S_i and S_{i+1} , then, the last cuts of e are $S_1 \cup \dots \cup S_i$ and $S_{i+1} \cup \dots \cup S_k$.

Note that the last cuts of a top edge are critical cuts, meanwhile, last cuts of a bottom edge are not necessarily a critical cut.

► **Definition 7** (Going Higher). We say an edge $e \in \delta(S)$ goes higher in S if the lowest critical cut S' such that $S \subsetneq S'$ satisfies $e \in \delta(S')$. Additionally, when S is a critical cut, by $\delta^\uparrow(S)$ we denote the edges in $\delta(S)$ that go higher in S . Similarly, $\delta^\rightarrow(S) = \delta(S) \setminus \delta^\uparrow(S)$ denotes the edges in $\delta(S)$ that don't go higher.



(a) C is a **degree cut** and C_i are **top cuts**. The black edges are **top edges**. e goes higher in C_1 , while a does not. The **last cuts** of a are C_1 and C_2 .



(b) S is a **cycle cut** and S_i are **bottom cuts**. The black edges are **bottom edges**. g and h go higher in S_1 , while c and d do not. c and d are **companions** and the **last cuts** of them are S_1 and $S_2 \cup S_3$. Moreover, g, h are **cycle partners**.

3.2 Constructing the O -Join Solution

Given a tree T sampled from the max-entropy distribution, we will describe a randomized process to construct a feasible solution for the O -Join formulation (Equation (3)) where O is the odd degree vertices of T .

Before describing the construction, we will restate the definition of even at last edges from [22].

► **Definition 8** (Definition 4.3 in [22]). For an edge e we say e is **even at last** if the two last cuts of e are even. If e is a bottom edge, this is equivalent to defining e to be even at last when all the min cuts containing e on the cycle defined by the graph consisting of S_e with $V \setminus S_e$ contracted are even. If e is a top edge, then it is even at last if its last cuts are even simultaneously.

Let x be the optimal solution of the subtour LP (Equation (1)). We will initialize $y = \frac{x}{2}$ so that y satisfies the O -Join constraints. Now, when an edge e is even at last, we will decrease y_e by τ , where τ is a parameter we will set later. Since an even at last edge can still cover lower level min-cuts in the hierarchy that are odd in T , we will increase the value of y_e accordingly to make y a feasible O -Join solution.

We utilize the fact that when an edge e is even at last, lower level cuts S such that $e \in \delta(S)$ are (in most cases) still even with probability $\Omega(1)$. For an edge e , let p_e denote the probability of e being even at last. Unfortunately, some edges can have $p_e \approx 0$ (see [23]). This is an issue for arguments that bound the contribution of each edge individually. Therefore, deviating from prior work, we will instead look at the expected number of even at last edges in $\delta(S)$ for a min cut S . This value is denoted by $p(\delta(S)) = \sum_{e \in \delta(S)} p_e$. We show lower bounds for this value for every min cut S , which in turn gives that $y(\delta(S))$ decreases meaningfully for each min cut S in the O -Join solution we construct. Moreover, when we increase edges to cover the violated O -Join constraint of a cut, we increase them according to (roughly) their p_e value. This in turn means that edges that increase in the third step of our construction, should also decrease meaningfully in the second step. A more accurate and complete description of this process is provided in Section 5.

Our main goal is to show that the O -Join solution of a tree sampled from the max entropy algorithm has cost at most $0.49776 \cdot c(x)$ in expectation. This immediately proves Theorem 1 as the O -Join polyhedron (Equation (3)) has an integrality gap of 1. To do so, we will show the O -Join solution y has expected cost at most $0.49776 \cdot c(x)$. Instead of bounding the contribution of each edge, we will bound the expectation on each minimum cut S as follows:

► **Lemma 9.** *There exists a randomized O -Join solution y for the random tree T sampled from the max entropy distribution such that for each min cut S we have,*

$$\mathbb{E}[y(\delta(S))] \leq 1 - 0.00448 = 0.99552$$

To analyze the cost of our solution, we utilize the dual formulation of the subtour LP (Equation (2)). Now, we will use Lemma 9 to prove Theorem 1.

► **Theorem 1.** *Given an optimal solution x to the subtour LP with $x_e \in \{0, \frac{1}{2}, 1\}$ for all $e \in E$, the max entropy algorithm produces a solution of cost at most $1.49776 \cdot c(x)$ in expectation.*

Proof. By complementary slackness we know for an edge e in the support of x , $c_e = \sum_{S: e \in \delta(S)} z_S$, therefore, the cost of the O -Join solution y can be written as,

$$c(y) = \sum_e c_e y_e = \sum_e \sum_{S: e \in \delta(S)} z_S y_e = \sum_{S: e \in \delta(S)} z_S \cdot y(\delta(S))$$

Now, by Lemma 9,

$$\mathbb{E}[c(y)] \leq \sum_{S: e \in \delta(S)} 0.99552 \cdot z_S = 0.49776 \cdot c(x)$$

where the final equality follows from strong duality. ◀

To prove Lemma 9, we analyze top cuts and bottom cuts separately. For each cut S , we show that either the value of y_e decreases for every edge in $\delta(S)$, or, if there exists an edge for which y_e does not decrease significantly, then the remaining edges in $\delta(S)$ have a larger decrease.

4 Probabilistic and Structural Lemmas

In this section, we will provide some key probabilistic and structural lemmas on the max entropy algorithm. These lemmas will provide strong probabilistic bounds as well as crucial observations about the structure of critical cuts that are essential in proving Lemma 9.

Before doing so, in Section 4.1 and Section 4.2 we introduce some key additional preliminaries that were omitted in Section 2.

4.1 Strongly Rayleigh Distributions

Given a distribution $\mu : \mathbb{Z}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$ over ground set $[n]$, its generating polynomial g_μ is defined

$$g_\mu(z) = \sum_{\kappa \in \mathbb{Z}_{\geq 0}^n} \mu(\kappa) z^\kappa$$

where $z^\kappa = \prod_{i=1}^n z_i^{\kappa_i}$. μ is **strongly Rayleigh** (SR) [2] if $g_\mu(z)$ is real stable. A polynomial $p(z)$ is real stable if $p(z) \neq 0$ for all $z \in \mathbb{C}^n$ with $\text{Im}(z_i) > 0$ for all $i \in [n]$. In other words, p is strongly Rayleigh if it has no zeros in the upper half of the complex plane. λ -uniform spanning tree distributions are strongly Rayleigh (see e.g. [2, 29]).

Negative Association. SR distributions are negatively associated [2, 9]. In particular, given any increasing functions $f, g : 2^E \rightarrow \mathbb{R}$ that depend on disjoint coordinates:

$$\mathbb{E}_\mu[f] \cdot \mathbb{E}_\mu[g] \geq \mathbb{E}_\mu[f \cdot g]$$

An easy consequence is the following:

► **Fact 10** (Fact 3.16 in [22]). *For any λ -uniform spanning tree distribution μ , for any $S \subseteq E$, any $k \geq \mathbb{R}$, and any $e \in E$ we have:*

1. *If $e \notin S$ then $\mathbb{P}_\mu[e \in T \mid S_T \geq k] \leq \mathbb{P}_\mu[e \in T]$ and $\mathbb{P}_\mu[e \in T \mid S_T \leq k] \geq \mathbb{P}_\mu[e \in T]$.*
2. *If $e \in S$, then $\mathbb{P}_\mu[e \in T \mid S_T \geq k] \geq \mathbb{P}_\mu[e \in T]$ and $\mathbb{P}_\mu[e \in T \mid S_T \leq k] \leq \mathbb{P}_\mu[e \in T]$.*

There is also a useful consequence of negative association when applied to a homogeneous distribution. (Recall that a polynomial is homogeneous when all terms have the same degree. A distribution is homogeneous when all outcomes have the same number of elements.)

► **Fact 11** (Fact 3.17 in [22]). *For any λ -uniform spanning tree distribution μ , for any set of edges $S \subseteq E$ and any $e \notin S$, we have:*

$$\mathbb{E}_{T \sim \mu}[S_T] \leq \mathbb{E}_{T \sim \mu}[S_T \mid e \notin T] \leq \mathbb{E}_{T \sim \mu}[S_T] + x_e$$

and similarly,

$$\mathbb{E}_{T \sim \mu}[S_T] - x_e \leq \mathbb{E}_{T \sim \mu}[S_T \mid e \in T] \leq \mathbb{E}_{T \sim \mu}[S_T]$$

When we apply one of these two facts, we will often simply say we are using negative association.

Closure Properties. A second consequence of real stability is that given an SR distribution μ , the following distributions are SR as well:

- **Projection.** $\mu|_S$, the projection of μ to the coordinates in some $S \subseteq [n]$.
- **Conditioning on a binary element to be 0 or 1.** If $z_i \in \{0, 1\}$, then $\mu|_{z_i=0}$ and $\mu|_{z_i=1}$ are SR.

Hoeffding's Theorem. For any subset $S \subseteq E$, the law of S_T for $T \sim \mu$ is distributed as the sum of independent Bernoulli random variables (not necessarily all with the same success probability). This is a consequence of the fact that the coefficients of any real rooted polynomial with positive coefficients can be described by a sum of Bernoullis [30, 2]. This makes the law of S_T particularly easy to analyze for any $S \subseteq E$, especially when one applies the following theorem of Hoeffding:

► **Theorem 12** ([20, Corollary 2.1]). Let $g : \{0, 1, \dots, n\} \rightarrow \mathbb{R}$ and $0 \leq q \leq n$ for some integer $n \geq 0$. Let $B_1, \dots, B_n$ be n independent Bernoulli random variables with success probabilities $p_1, \dots, p_n$, where $\sum_{i=1}^n p_i = q$ that minimizes (or maximizes)

$$\mathbb{E}[g(B_1 + \dots + B_n)]$$

over all such distributions. Then, $p_1, \dots, p_n \in \{0, x, 1\}$ for some $0 < x < 1$. In particular, if only m of p_i 's are nonzero and ℓ of p_i 's are 1, then the remaining $m - \ell$ are $\frac{q-\ell}{m-\ell}$.

A very useful corollary is the following.

► **Lemma 13** (Lemma 3.23 of [22]). Let $S \subseteq V$ with $x(\delta(S)) = 2$ and $|\delta(S)| \leq 4$. Then $\mathbb{P}[\delta_T(S) \text{ even}] \geq 13/27$.

The proof is omitted in [22] as it follows straightforwardly from Theorem 12. One can see Lemma 3.6 in [26] for a proof.

Finally, we will need the following lemma.

► **Lemma 14** (Lemma 3.21 of [22]). Let $S \subseteq E$ with $|S| = 3$. Furthermore, assume that $\mathbb{P}[|S \cap T| \geq 1] = 1$. Then, $\mathbb{P}[|S \cap T| = 1] \geq \frac{1}{2}$ and $\mathbb{P}[|S \cap T| = 2] \geq \frac{3}{8}$.

4.2 Polynomial Capacity

The *capacity* at $\alpha \in \mathbb{N}_+^n$ of a real stable polynomial $p(x_1, \dots, x_n)$ with positive coefficients is defined as:

$$\text{cap}_\alpha(p) = \inf_{x \in \mathbb{R}_{>0}^n} \frac{p(x)}{x^\alpha}$$

A classical result of Gurvits [14] relates the capacity of a polynomial to the coefficient of $\prod_{i=1}^n x_i$ for n -variate homogeneous polynomials of degree n as follows (where $\mathbf{1}$ is the vector consisting of n 1s):

► **Theorem 15** ([15]). Let $p(x_1, \dots, x_n)$ be a homogeneous real stable polynomial of degree n with non-negative coefficients. Then, where $C_i = \min(\deg_p(i), i)$,

$$\frac{\partial^n}{\partial x_1 \dots \partial x_n} p \Big|_{x_1=\dots=x_n=0} \geq \text{cap}_{\mathbf{1}}(p) \prod_{i=2}^n \left(\frac{C_i - 1}{C_i} \right)^{C_i - 1}$$

Note that $\frac{\partial^n}{\partial x_1 \dots \partial x_n} p \Big|_{x_1=\dots=x_n=0}$ is exactly the coefficient of $\prod_{i=1}^n x_i$. There are various similar statements in the literature, and we will use the following, first stated as Theorem 5.1 of [15] and restated as follows in [4]:

► **Theorem 16** ([15, 4]). Let p be a homogeneous real stable polynomial of degree d with positive coefficients. Let $\alpha \in \mathbb{N}_+^n$ such that $\sum_{i=1}^n \alpha_i = d$. For $i < n$, let d_i be the degree of x_i in the polynomial

$$\partial_{i+1}^{\alpha_{i+1}} \dots \partial_n^{\alpha_n} \Big|_{x_{i+1}=\dots=x_n=0}$$

and d_n the degree of x_n in p . Then, where p_α is the coefficient of the term $\prod_{i=1}^n x_i^{\alpha_i}$,

$$\frac{p_\alpha}{\text{cap}_\alpha(p)} \geq \prod_{i=2}^n \binom{d_i}{\alpha_i} \frac{\alpha_i^{\alpha_i} (d_i - \alpha_i)^{d_i - \alpha_i}}{d_i^{d_i}}$$

Furthermore, the capacity of a real stable polynomial can be bounded using its gradient. In particular, we can apply the following theorem of Gurvits and Leake [17] (also see [16]) generalizing [13]. (We do not need the generalization here, but we state the stronger form regardless.)

► **Theorem 17** ([17]). *Let p be a real stable polynomial in n variables with non-negative coefficients, and fix any $\alpha \in \mathbb{N}_+^n$. If $p(\mathbf{1}) = 1$ and $\|\alpha - \nabla p(\mathbf{1})\| < 1$, then*

$$\text{cap}_\alpha \geq (1 - \|\alpha - \nabla p(\mathbf{1})\|)^n$$

We will use the following corollary in this work, which follows easily from the above.

► **Corollary 18.** *Let $p(x_1, \dots, x_n)$ be the generating polynomial of a strongly Rayleigh distribution μ over ground set $e_1, \dots, e_n$. If $\mathbb{E}[e_i] = 1$ for all $1 \leq i \leq n$, or equivalently $\nabla p(\mathbf{1}) = \mathbf{1}$, then if d_i is the maximum degree of x_i ,*

$$p_{\mathbf{1}} \geq \prod_{i=1}^n \frac{d_i(d_i - 1)^{d_i - 1}}{d_i^{d_i}}$$

Proof. Let d be the maximum degree of p . Let $p^H(z, x_1, \dots, x_n) = z^d(x_1/z, \dots, x_n/d)$ be the homogenization of p . Then, $\nabla p^H(\mathbf{1}) = (d - n, 1, \dots, 1)$. Set $\alpha = (d - n, 1, \dots, 1)$. By Theorem 17, $\text{cap}_{(d-n, 1, \dots, 1)} = 1$. Applying Theorem 16 with α , noting $\alpha_i = 1$ for $i \geq 2$, we obtain:

$$p_\alpha^H \geq \prod_{i=1}^n \frac{d_i(d_i - 1)^{d_i - 1}}{d_i^{d_i}}$$

where we adjust the indices to match the degree of each variable x_i . But $p_\alpha^H = p_{\mathbf{1}}$, so the corollary follows. ◀

4.3 Structure of Critical Cuts

In this section, we recall some basic facts about the structure of critical cuts. For proofs, we refer the interested reader to Section 3 of [22]. These facts are used solely to ensure that our case analysis is exhaustive and covers all possible scenarios.

► **Fact 19** (Fact 3.10 [22]). *Suppose that S is a critical set. If some (contracted) vertex $v \in S$ has two edges to $w := V \setminus S$, then S is a cycle cut.*

► **Fact 20** (Fact 3.11 [22]). *Suppose that S and S' are two distinct tight sets. Then $|\delta(S) \cap \delta(S')| \leq 2$.*

► **Fact 21** (Fact 3.12 [22]). *Suppose that S and S' are two critical sets such that $S \subset S'$. Then if $|\delta(S) \cap \delta(S')| = 2$, then S' is a cycle cut.*

► **Fact 22** (Fact 3.13 [22]). *Suppose that $S \subset S'$ are two critical cycle cuts. Then any two edges are cycle partners on at most one of these (cycle) cuts.*

The following two corollaries are immediate.

► **Corollary 23.** *Suppose S is a degree cut. Then S has at most one edge that goes higher.*

► **Corollary 24.** *Suppose S is a cycle cut. Then S has either exactly two edges or no edges that go higher.*

We will also prove the following simple fact.

► **Fact 25.** *Let S be a min cut that is not the last cut for any edge in the support graph. S is always even in the tree T .*

Proof. Let S' be the parent of S in the hierarchy of critical cuts. By Fact 4, S' must be a cycle cut with child cuts $S_1, \dots, S_k$ with two edges between S_i and S_{i+1} for $1 \leq i \leq k-1$. Since S is not the last cut of any edge, it must be of the form $S_i \cup \dots \cup S_j$ for $1 < i < j < k$.

Since exactly one of the two companions between S_{i-1}, S_i and S_j, S_{j+1} are in T , $|\delta_T(S)| = 2$ and S is even. ◀

4.4 Bounding the Expected Number of Even at Last Edges on Min Cuts

As previously mentioned, for an edge e , there cannot be any meaningful lower bound on the probability of e being even at last as there can be edges with $p_e \approx 0$. In contrast, we will show that for any top cut S , there are strong lower bounds on $p(\delta^{\rightarrow}(S))$. This intuitively shows that in expectation, $y(\delta(S))$ can be decreased for every min cut.

The following lemma can be thought of as the analog of Lemma 5.3 of [22] (which showed at least one edge in each cut has $p_e \geq \frac{1}{27}$) when adapted to our framework, and uses similar proof ideas.

► **Lemma 26.** *For any top cut S with no edge going higher, $p(\delta(S)) \geq \frac{4}{27}$. Moreover, if $W \subset \delta(S)$, $|W| = 3$. Then, $p(W) \geq \frac{1}{27}$.*

Proof. Let v be the vertex corresponding to S after contraction. For an edge $e \in \delta(S)$, with (contracted) last cuts u and v , we have,

$$\begin{aligned} \mathbb{P}[e \text{ Even at last}] &= 1 - \mathbb{P}[\delta_T(v) \text{ is odd} \vee \delta_T(u) \text{ is odd}] \\ &= 1 - \mathbb{P}[\delta_T(u) \text{ is odd}] - \mathbb{P}[\delta_T(v) \text{ is odd}] + \mathbb{P}[\delta_T(v) \text{ is odd} \wedge \delta_T(u) \text{ is odd}] \\ &\geq \frac{13}{27} - \mathbb{P}[\delta_T(v) \text{ is odd}] + \mathbb{P}[\delta_T(v) \text{ is odd} \wedge \delta_T(u) \text{ is odd}] \end{aligned}$$

Where in the last line we used Lemma 13. Now we will bound $\mathbb{P}[\delta_T(v) \text{ is odd} \wedge \delta_T(u) \text{ is odd}]$.

$$\begin{aligned} \mathbb{P}[\delta_T(v) \text{ is odd} \wedge \delta_T(u) \text{ is odd}] &= \mathbb{P}[\delta_T(u) \text{ is odd} \mid \delta_T(v) \text{ is odd}] \cdot \mathbb{P}[\delta_T(v) \text{ is odd}] \\ &\geq \mathbb{P}[\delta_T(u) \text{ is odd} \mid \delta_T(v) = 1] \cdot \mathbb{P}[\delta_T(v) = 1] \end{aligned}$$

Conditioning on $\delta_T(v) = 1$ is identical to conditioning on $\delta_T(v) \leq 1$, as $\mathbb{P}[\delta_T(v) \geq 1] = 1$. Therefore, by negative association, we have $\mathbb{P}[e \in T \mid \delta_T(v) = 1] \leq \frac{1}{2}$. Moreover,

$$\mathbb{P}[\delta_T(u) \text{ is odd} \mid \delta_T(v) = 1] \geq \mathbb{P}[e \in T \mid \delta_T(v) = 1]$$

as the measure conditioned on $\delta_T(v) = 1$ first samples a tree T' from $V - v$ and then chooses an edge in $\delta_T(v)$ according to its conditional probability. Depending on the degree of u in T' being even or odd, we can make $\delta_T(u)$ odd by either conditioning on e being in the tree or out of the tree. By negative association, $\mathbb{P}[e \notin T \mid \delta_T(v) = 1] \geq \mathbb{P}[e \in T \mid \delta_T(v) = 1]$, therefore,

$$\mathbb{P}[\delta_T(v) \text{ is odd} \wedge \delta_T(u) \text{ is odd}] \geq \mathbb{P}[e \in T \mid \delta_T(v) = 1] \cdot \mathbb{P}[\delta_T(v) = 1]$$

Which in turn gives,

$$p_e \geq \frac{13}{27} - \mathbb{P}[\delta_T(v) \text{ is odd}] + \mathbb{P}[e \in T \mid \delta_T(v) = 1] \cdot \mathbb{P}[\delta_T(v) = 1]$$

As $\sum_{e \in \delta(v)} \mathbb{P}[e \in T \mid \delta_T(v) = 1] = 1$, summing over $\delta(v)$ gives,

$$\begin{aligned} p(\delta(v)) &\geq \frac{52}{27} - 4 \cdot \mathbb{P}[\delta_T(v) \text{ is odd}] + \mathbb{P}[\delta_T(v) = 1] \\ &= \frac{52}{27} - 3 \cdot \mathbb{P}[\delta_T(v) = 1] - 4 \cdot \mathbb{P}[\delta_T(v) = 3] \end{aligned}$$

By Theorem 12, this term attains its minimum when the Bernoullis are $\{1, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$ at $\frac{4}{27}$. This shows the first claim in the lemma, that $p(\delta(S)) \geq \frac{4}{27}$.

For the second part, consider the inequality,

$$p_e \geq \frac{13}{27} - \mathbb{P}[\delta_T(S) \text{ is odd}] + \mathbb{P}[e \in T \mid \delta_T(v) = 1] \cdot \mathbb{P}[\delta_T(v) = 1]$$

Since for every edge in $\delta(S)$ we have $\mathbb{P}[e \in T \mid \delta_T(v) = 1] \leq \frac{1}{2}$, $\mathbb{E}[W_T \mid \delta_T(v) = 1] \geq \frac{1}{2}$. By summing this inequality over all $e \in W$ we get,

$$\begin{aligned} p(W) &\geq \frac{39}{27} - 3 \cdot \mathbb{P}[\delta_T(v) \text{ is odd}] + 0.5 \cdot \mathbb{P}[\delta_T(v) = 1] \\ &= \frac{13}{9} - 2.5 \cdot \mathbb{P}[\delta_T(v) = 1] - 3 \cdot \mathbb{P}[\delta_T(v) = 3] \end{aligned}$$

By Theorem 12, this term attains its minimum at $\frac{1}{27}$ which concludes the proof. $\blacktriangleleft$

Now, we will prove a similar result for top cuts with an edge that goes higher. But first, we will prove the following simple lemma.

► **Lemma 27.** *Let e be a top edge with (contracted) last cuts u and v . Suppose u has an edge going higher, and v does not. Then e is even at last with probability at least $\frac{13}{54}$.*

Proof. By Lemma 13, v is even at last with probability at least $\frac{13}{27}$. Additionally, as e is in a different level of the hierarchy than the edges in $\delta^{\rightarrow}(v)$, by letting e be in or out of T respectively, we can fix the parity of u with probability $\frac{1}{2}$. This concludes the proof. $\blacktriangleleft$


► **Lemma 28.** *Let S be a top cut with an edge $e \in \delta(S)$ that goes higher and let W be any two of the edges in $\delta^{\rightarrow}(S)$. Then, $p(W) \geq \frac{7}{32}$.*

Proof. Denote the two edges in W by a, b . If the other last cut of any of these two doesn't go higher, by Lemma 27 $p(W) \geq \frac{13}{54}$ which satisfies the lower bound. Now, assume the other last cut of each of a and b has an edge that goes higher. Call these other last cuts by S_a, S_b and the edges going higher from them e_a, e_b . Now, condition on $e \in T$. Since e goes higher in S , it's independent from the edges that are inside S . Therefore, by Lemma 14, with probability at least $\frac{1}{2}$ exactly one of the edges in $\delta^{\rightarrow}(S)$ are in T . This makes the degree of S even.

Similarly, condition on $e \notin T$ which happens with probability $\frac{1}{2}$. By Lemma 14, with probability at least $\frac{3}{8}$ exactly 2 edges in $\delta^{\rightarrow}(S)$ are in T which makes S even. Thus,

$$\mathbb{P}[e \in T \wedge \delta_T(S) \text{ is even}] = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4} \quad \text{and} \quad \mathbb{P}[e \notin T \wedge \delta_T(S) \text{ is even}] = \frac{1}{2} \cdot \frac{3}{8} = \frac{3}{16}$$

where we have used the fact that e is independent from the edges in $\delta^{\rightarrow}(S)$.



■ **Figure 2** Illustration of Lemma 28 when S_a and S_b both have an edge going higher. The green edges represent the edges in W .

Now, in both cases, depending on the parity of S_a, S_b inside S , we can choose e_a (or e_b) to be inside or outside T to make a (or b) even at last. Since $\mathbb{P}[e_a \in T \mid e \in T] \leq \frac{1}{2} \leq \mathbb{P}[e_a \notin T \mid e \in T]$ and $\mathbb{P}[e_a \notin T \mid e \notin T] \leq \frac{1}{2} \leq \mathbb{P}[e_a \in T \mid e \notin T]$, we have,

$$\begin{aligned} p(W) &\geq \frac{1}{4} \left(\mathbb{P}[e_a \in T \mid e \in T] + \mathbb{P}[e_b \in T \mid e \in T] \right) + \frac{3}{16} \left(\mathbb{P}[e_a \notin T \mid e \notin T] + \mathbb{P}[e_b \notin T \mid e \notin T] \right) \\ &\geq \frac{7}{32} \end{aligned}$$

Where the last inequality follows from negative association. ◀

For bottom cuts we can prove stronger bounds. In fact, bottom edges inside a cycle cut are all simultaneously even at last. This symmetry enables us to individually bound p_e for each edge e .

► **Lemma 29.** *Every bottom edge is even at last with probability at least $\frac{1}{4}$.*


Proof. Consider a bottom edge e , so that S_e is a cycle cut with child cuts $S_1, \dots, S_k$ with two edges between each S_i and S_{i+1} for each $1 \leq i \leq k-1$. When a tree on S_e is chosen, we will obtain exactly one edge among every pair of adjacent child cuts. The edges in $\delta(S)$ are comprised of two sets of edges, $A = \{a, b\} = \delta(S) \cap \delta(S_1)$ and $B = \{c, d\} = \delta(S) \cap \delta(S_k)$. So, e is even at last exactly when $A_T = B_T = 1$.

Project μ to $\{a, b, c, d\}$. The resulting distribution has generating polynomial $p(x_a, x_b, x_c, x_d)$. Symmetrize so that $p'(x_a, x_b) = p(x_a, x_a, x_b, x_b)$. By Corollary 18, where d_1 is the degree of x_a and d_2 the degree of x_b ,

$$\mathbb{P}[A_T = B_T = 1] = p'_{(1,1)} \geq \frac{d_1(d_1 - 1)^{d_1 - 1}}{d_1^{d_1}} \cdot \frac{d_2(d_2 - 1)^{d_2 - 1}}{d_2^{d_2}} \geq \frac{1}{4}$$

as desired, since $A_T \leq 2$ and $B_T \leq 2$ so $d_1 \leq 2$ and $d_2 \leq 2$. ◀

We remark that this lemma is tight: there are instances where this probability is exactly $\frac{1}{4}$, and the relevant edges have generating polynomial $(x_a + x_c)(\frac{1}{2} + x_b)(\frac{1}{2} + x_d)$.



■ **Figure 3** The structure of S_e at the proof of Lemma 29.

4.5 Structure of Degree Cuts and K_5

In this section, we show that if a degree cut S satisfies a certain structure, then G_S must be the graph K_5 . We use this fact in Section 5 to show that if many edges in $\delta^{\rightarrow}(S)$ have conditions that prevent them from decreasing reasonably, the degree cut S must have a simple structure. In particular, G_S must be a K_5 . We note that Gupta et al. [12] also treated K_5 as a special case, and adapted the algorithm to treat these cuts differently. While we do not change the algorithm, we similarly analyze this case separately.

► **Lemma 30.** *Let S be a degree cut with (contracted) cuts $u, v, w \in G[S]$ each with an edge that goes higher. If, u, v, w forms a triangle, then G_S is K_5 .*

Proof. Consider the set $R = S \setminus \{u, v, w\}$. $|\delta(R)| = 4$, so R is a min-cut. However, $R \subsetneq S$, and since there are no proper min-cuts inside a degree-cut, it must be the case that $|R| = 1$. Therefore G_S is a 4-regular graph with 5 vertices which concludes the proof. ◀


The following results from the fact that setting $\lambda = 1$ is the λ -uniform distribution for K_4 when $x_e = \frac{1}{2}$ for all edges. (Note λ is unique for a fixed x , so this is the only possible distribution.)

► **Fact 31.** *The max entropy distribution on the graph K_4 is the uniform spanning tree distribution.*

We will now complement Lemma 30 with the following.

► **Lemma 32.** *Let S be a degree cut of the support graph G such that $G_S = K_5$. Every edge in $G[S]$ is even at last with probability at least $\frac{1}{4}$.*

We omit the simple proof in this extended abstract.



■ **Figure 4** A degree cut where G_S is a K_5 .

5 Analysis

Now, we describe the construction of the O-Join solution in full detail.

For an edge e , recall p_e is the probability of e being even at last. For top edges define $\tilde{p}_e = \min(\alpha, p_e)$ and for bottom edges $\tilde{p}_e = \min(\beta, p_e)$ for constants α, β we will set momentarily. Note that $\tilde{p}(A)$ for $A \subseteq E$ denotes $\sum_{e \in A} \tilde{p}_e$. For each edge e with $p_e \neq 0$, let B_e be an independent Bernoulli with success probability $\frac{\tilde{p}_e}{p_e}$, if $p_e = 0$, let $B_e = 0$. Moreover, for two bottom edges in the same cycle cut e, f , let $B_e = B_f$ with probability one. Note that since these bottom edges are always simultaneously even at last this is well-defined.

Finally, let S be a critical cut and $e \in \delta^{\rightarrow}(S)$. By $r_S(e)$ we denote $\frac{\tilde{p}_e}{\tilde{p}(\delta^{\rightarrow}(S))}$. In other words, $r_S(e)$ represents the fraction of even at last edges in $\delta^{\rightarrow}(S)$ that e is responsible for.

We will now construct y in 3 different steps, done sequentially. The construction is done after sampling a tree T and sampling the Bernoullis B_e for every $e \in E$.

1. Let $y_e = \frac{1}{4}$ for each $e \in E$.
2. For any even at last edge e with $B_e = 1$, reduce y_e by τ .
3. For each odd min cut S , let $\Delta(S) = \max(0, 1 - y(\delta(S)))$. For each edge e and its last cuts S, S' , increase e by $\max(r_S(e) \cdot \Delta(S), r_{S'}(e) \cdot \Delta(S'))$.

Throughout the analysis let $\beta = \frac{1}{4}, \alpha = 0.1129032, \tau = \frac{1}{12}$. The third step of our construction along with Fact 25, ensure that y satisfies the O-Join constraint for all odd min cuts. Moreover, since $\tau = \frac{1}{12}$, the O-Join constraint for every non min cut S is also satisfied, even if all edges in $\delta(S)$ are simultaneously reduced, as such cuts have at least 6 edges covering them.

Now, we will consider different cases for a min-cut S and show that for every case $\mathbb{E}[y(\delta(S))]$ is at least reduced by 0.00448 in expectation. For brevity, we do not consider cuts of the final cycle considered in Algorithm 1, as these cuts can easily be seen to obey the bounds described here.

► **Lemma 33.** *Let S be a top cut with an edge that goes higher, then, $\tilde{p}(\delta(S)) \geq 2\alpha$.*

Proof. Let a, b, c be the three edges in $\delta(S)$ that don't go higher. If the probability of being even at last on none of these edges is more than α , then,

$$\tilde{p}(\delta(S)) = p_a + p_b + p_c = \frac{(p_a + p_b) + (p_b + p_c) + (p_a + p_c)}{2} \stackrel{(i)}{\geq} \frac{3}{2} \cdot \frac{7}{32} \geq 2\alpha$$

where inequality (i) follows from Lemma 28. Otherwise, assume $\tilde{p}(a) = \alpha$, then, by Lemma 28, $\tilde{p}(b) + \tilde{p}(c) \geq \min(\alpha, \frac{7}{32}) = \alpha$. Therefore, in all cases $\tilde{p}(\delta(S)) \geq 2\alpha$. ◀

A crucial consequence of Lemma 33 is that for any top cut S and an edge $a \in \delta^{\rightarrow}(S)$, we have $r_S(a) \leq \frac{1}{2}$ as $\tilde{p}_a \leq \alpha$ by definition.

► **Lemma 34.** *Let S be a top cut with no edge that goes higher, then,*

$$\mathbb{E}[y(\delta(S))] \leq 1 - \left(\alpha + \frac{1}{27} - \frac{\beta}{4} \right) \cdot \tau$$

Proof. Let a, b, c, d be the edges in $\delta^{\rightarrow}(S)$ and denote the other last cuts of these edges respectively by S_a, S_b, S_c, S_d . For an edge a , if S_a has an edge that goes higher, then by Lemma 27, $\tilde{p}_a = \alpha$. Let e_a be the edge in $\delta(S_a)$ that goes higher.

If e_a is a bottom edge, then

$$\tilde{p}(e_a) = \beta \quad \text{and} \quad \mathbb{P}[S_a \text{ is odd} \mid a \text{ is reduced}] = \frac{1}{2}$$

since by switching e_a with its companion, we can change the parity of S_a while e_a remains even at last.

To cover for the deficit caused by e_a , a increases by at most $\frac{\beta}{4} \cdot \tau$ in expectation. This is because by Lemma 33, a is at most responsible for half of $\tilde{p}(\delta^{\rightarrow}(S_a))$.

Otherwise, if e_a is a top edge, by Lemma 14 and our construction of the O-Join solution,

$$\tilde{p}(e_a) = \alpha \quad \text{and} \quad \mathbb{P}[S_a \text{ is odd} \mid a \text{ is reduced}] \leq \frac{5}{8}$$

Consequently, a increases by at most $\frac{5\alpha}{16} \cdot \tau \leq \frac{\beta}{4} \cdot \tau$. Therefore, for any edge a such that S_a has an edge that goes higher, $\mathbb{E}[y_a] \leq \frac{1}{4} - (\alpha - \frac{\beta}{4}) \cdot \tau$.

■ **Case 1:** Two or more of S_a, S_b, S_c, S_d have an edge that goes higher. As $\alpha - \frac{\beta}{4} > 0$,

$$\mathbb{E}[y(\delta(S))] \leq 1 - 2 \cdot \left(\alpha - \frac{\beta}{4} \right) \cdot \tau$$


■ **Case 2:** Only S_a has an edge that goes higher (see Figure 5). As above, $\mathbb{E}[y_a] \leq \frac{1}{4} - (\alpha - \frac{\beta}{4}) \cdot \tau$. For the remaining three edges, by Lemma 26, $\tilde{p}(\{b, c, d\}) \geq \frac{1}{27}$. Therefore,

$$\mathbb{E}[y(\delta(S))] \leq 1 - \left(\alpha + \frac{1}{27} - \frac{\beta}{4} \right) \cdot \tau$$

■ **Case 3:** None of S_a, S_b, S_c, S_d has an edge that goes higher. Since $\alpha \leq \frac{4}{27}$, by Lemma 26,

$$\mathbb{E}[y(\delta(S))] \leq 1 - \alpha \cdot \tau$$

Hence, for every case, $\mathbb{E}[y(\delta(S))] \leq 1 - \left(\alpha + \frac{1}{27} - \frac{\beta}{4} \right) \cdot \tau$ ◀



■ **Figure 5** Illustration of the worst case in Lemma 34. S' is a degree cut and e is a bottom edge.

► **Lemma 35.** Let e be a bottom edge inside a cycle cut S with child cuts $S_1, \dots, S_k$ with two edges between each S_i and S_{i+1} for each $1 \leq i \leq k-1$. Also, let S' be the parent of S in the cut hierarchy. Then, we will bound the amount e increases to cover the deficit on its last cuts caused by decreasing edges in $\delta(S)$ in three cases:

1. If S' is a degree cut with no edge going higher, e increases by at most $2\alpha \cdot \tau$.
2. If S' is a degree cut with an edge going higher, e increases by at most $\left(\frac{3\alpha}{2} + \frac{\beta}{4} \right) \cdot \tau$.
3. If S' is a cycle cut, e increases by at most $\left(\frac{3\beta}{4} \right) \cdot \tau$.

By Corollaries 23 and 24, the structure of S' falls into one of these three categories.

We omit the proof in this extended abstract.

21:18 Dual Charging for Half-Integral TSP


► **Lemma 36.** *Let S be a bottom cut with no edge that goes higher, then,*

$$\mathbb{E}[y(\delta(S))] \leq 1 - (3\beta - 6\alpha) \cdot \tau$$

Proof. By Lemma 29, for every bottom edge e , $\tilde{p}(e) = \frac{1}{4} = \beta$. Additionally, by Lemma 35 for any $e \in \delta(S)$,

$$\mathbb{E}[y(e)] \leq \frac{1}{4} - \left(\beta - \max \left(2\alpha, \frac{3\alpha}{2} + \frac{\beta}{4}, \frac{3\beta}{4} \right) \right) \cdot \tau = \frac{1}{4} - \left(\frac{3\beta}{4} - \frac{3\alpha}{2} \right) \cdot \tau.$$

Therefore, $\mathbb{E}[y(\delta(S))] \leq 1 - (3\beta - 6\alpha) \cdot \tau$. ◀



■ **Figure 6** Illustration of the worst case in Lemma 36. The edges f, h, g are top edges while e is a bottom edge.

Now that we have dealt with min cuts having no edges going higher, we now turn to those with at least one such edge. The structure of these cuts is known from Corollaries 23 and 24.

► **Lemma 37.** *Let S be a top cut with an edge e that goes higher. Let a, b, c be the edges in $\delta^{\rightarrow}(S)$, and let S_a, S_b , and S_c be the other last cuts of a, b , and c , respectively. If each of S_a, S_b , and S_c have an edge that goes higher, then $r_{S_a}(a) \leq \frac{1}{3}$. More generally, $\tilde{p}(\delta^{\rightarrow}(S_a)) = 2\alpha + \tilde{p}_a$*

Proof. Consider the min-cut S_a , and let $g \in \delta^{\rightarrow}(S_a) \setminus \{a\}$. Let R be the other last cuts of g . Suppose R has an edge that goes higher. Since at most four (contracted) min-cuts can have such an edge inside a degree cut, R must be either S_b or S_c .

In this case, the cuts S, S_a , and R form a triangle. By Lemma 30, G_S forms a K_5 , and Lemma 32 implies that g (and every edge in $\delta^{\rightarrow}(S_a)$) is even at last with probability $\frac{1}{4} \geq \alpha$. Therefore, $r_{S_a}(a) = \frac{1}{3}$.


Now, suppose that the other last cut of every edge in $\delta^{\rightarrow}(S_a) \setminus \{a\}$ does not have an edge that goes higher. In this case, by Lemma 27, these edges are even at last with probability at least $\frac{13}{54} \geq \alpha$. Thus concluding $r_{S_a}(a) \leq \frac{1}{3}$. Moreover, in both cases it is clear that $\tilde{p}(\delta^{\rightarrow}(S_a)) = 2\alpha + \tilde{p}_a$. ◀

We will utilize Lemma 37 to show $\mathbb{E}[y(\delta(S))]$ decreases meaningfully when S is a top cut with an edge going higher.

► **Lemma 38.** *Let S be a top cut with an edge that goes higher, then,*

$$\mathbb{E}[y(\delta(S))] \leq 1 - \left(\frac{7\alpha}{24} + \frac{\beta}{12} \right) \cdot \tau$$

We omit the proof in this extended abstract, which has extended casework.



■ **Figure 7** A min cut satisfying the conditions of Lemma 37.

We have proved upper bounds on $\mathbb{E}[y(\delta(S))]$ for all top cuts and all bottom cuts with no edge going higher. We next handle bottom cuts with edges going higher. By Corollary 24, such cuts have exactly two edges going higher.

► **Lemma 39.** *Let S be a bottom cut with two edges that go higher, then,*

$$\mathbb{E}[y(\delta(S))] \leq 1 - \left(\frac{2\beta}{3} - \alpha \right) \cdot \tau$$

Proof. Let S' be the parent of S in the hierarchy, by e, f denote the two edges in S that go higher, denote the other edges in $\delta(S')$ by g, h . Additionally, let C be the critical cut inside S' which g, h go higher from and let a, b be the edges in $\delta^{\rightarrow}(S)$. We will now carefully analyze $\mathbb{E}[y(\delta(S))]$.

■ **Case 1:** S' is a bottom cut. By Corollary 24, at least two of the edge in $\delta(S')$ are companions that don't go higher in S' , WLOG, let e, g be these edges. Since they are companions, e, g are even at last simultaneously. Moreover, if e is even at last, then, S' is even. This implies the parity of C and S should be the same and if S is odd, C should be odd too. Therefore, increasing a, b simultaneously covers the deficit caused by reducing e, g .

Another key observation is that a, b in total increase exactly as much as e has decreased. This implies that when e (and simultaneously g) is even at last and S is odd, $y(\delta(S))$ remains the same and we don't need to consider this scenario. By Lemma 35, e at most increases by $\left(\frac{3\alpha}{2} + \frac{\beta}{4} \right) \cdot \tau$. Meanwhile, since $\mathbb{P}[S \text{ is even} \mid e \text{ is reduced}] = \frac{1}{2}$, it can also decrease $y(\delta(S))$ by $\frac{\beta}{2} \cdot \tau$ on expectation. The same bound holds for f if it is a bottom edge. Now, assume f is a top edge and let c, d be the edges going higher from the last cuts of f . By Lemma 33, f can at most contribute to half of the even at last edges in its last cuts. Therefore, in expectation, f can at most increase by $\frac{\beta}{4}$ for each of its last cuts. Recall that one can swap c , or d , with their companion to fix the parity of the last cuts of f while retaining the even at last property for c , or d . Since $\frac{\beta}{2} > \frac{3\alpha}{2} - \frac{\beta}{4}$, the latter case is worse.

Finally, consider h , if h is a bottom edge, $\tilde{p}_h = \beta$. Moreover, $\mathbb{P}[S \text{ is odd} \mid h \text{ is reduced}] = \frac{1}{2}$. Thus, a, b each increase by $\frac{\beta}{4}$ in step (3) of our construction because of h . (Remember that $r_S(a), r_S(b) = \frac{1}{2}$.) If h was a top edge, then $\tilde{p}_h = \alpha$. Since $\alpha \leq \frac{\beta}{2}$, in the worst case h is a bottom edge. Using $\tilde{p}_a, \tilde{p}_b = \beta$ and all the preceding observations gives us,

$$\mathbb{E}[y(\delta(S))] \leq 1 - \left(2\beta + \frac{\beta}{2} - \left(\frac{3\alpha}{2} + \frac{\beta}{4} \right) - \frac{\beta}{2} - \frac{\beta}{2} \right) \cdot \tau = 1 - \left(\frac{5\beta}{4} - \frac{3\alpha}{2} \right) \cdot \tau$$

- **Case 2:** S' is a top cut. Since S' is a top cut, by Corollary 23, at least three of e, f, g, h are top edges in $\delta^{\rightarrow}(S)$. First, suppose e, f, g are those edges. Because of S' , e (and similarly f) needs to increase by at most $\frac{\beta}{2} \cdot r_S(e) \cdot \tau = \frac{\beta}{2} \cdot \frac{\tilde{p}_e}{\tilde{p}_e + \tilde{p}_f + \tilde{p}_g} \cdot \tau$. For its other last cut, it increases by at most $\frac{\beta}{2} \cdot \frac{\tilde{p}_e}{2} \cdot \tau$ by Lemma 33. Finally, g can increase a, b in total by at most $\tilde{p}_g$ in expectation. $\mathbb{E}[y(\delta(S))]$ is maximized when $\tilde{p}_e = \tilde{p}_f = \tilde{p}_g = \alpha$ giving

$$\mathbb{E}[y(\delta(S))] \leq 1 - \left(2\beta - 2 \cdot \left(\frac{\beta}{4} + \frac{\beta}{6} \right) - \alpha - \frac{\beta}{2} \right) \cdot \tau = 1 - \left(\frac{2\beta}{3} - \alpha \right) \cdot \tau$$

Otherwise, WLOG, e, g, h are top edges. Similar to the previous case, $\mathbb{E}[y(\delta(S))]$ is maximized when $\tilde{p}_e = \tilde{p}_f = \tilde{p}_g = \alpha$. If f is a bottom edge, similar to previous cases, it can decrease $\mathbb{E}[y(\delta(S))]$ by $\left(\frac{3\alpha}{2} - \frac{\beta}{4} \right) \cdot \tau$, so:

$$\mathbb{E}[y(\delta(S))] \leq 1 - \left(2\beta - \left(\frac{3\alpha}{2} - \frac{\beta}{4} \right) - \frac{\beta}{4} - \frac{\beta}{6} - \alpha - \alpha \right) \cdot \tau = 1 - \left(\frac{11\beta}{6} - \frac{7\alpha}{2} \right) \cdot \tau$$

However, if f is a top edge it can decrease $\mathbb{E}[y(\delta(S))]$ by at most $\frac{\beta}{2} \cdot \tau$. Meanwhile, the edge that goes higher from S' (if such an edge exists) is a top edge, therefore, to cover for the deficit on S' , e needs to increase by at most $\frac{5}{8} \cdot \alpha \cdot \tau$ (where we have used Lemma 14 to show $\mathbb{P}[S \text{ is odd} \mid f \text{ is reduced}] \leq \frac{5}{8}$). This gives

$$\mathbb{E}[y(\delta(S))] \leq 1 - \left(2\beta - \left(\frac{\beta}{4} + \frac{5\alpha}{24} \right) - \frac{\beta}{2} - \alpha - \alpha \right) \cdot \tau = 1 - \left(\frac{5\beta}{4} - \frac{53\alpha}{24} \right) \cdot \tau$$

as desired. ◀

Now, we are ready to prove the main lemma of the paper.

- **Lemma 9.** *There exists a randomized O-Join solution y for the random tree T sampled from the max entropy distribution such that for each min cut S we have,*

$$\mathbb{E}[y(\delta(S))] \leq 1 - 0.00448 = 0.99552$$

Proof. By Corollaries 23 and 24, every min-cut falls into one of the Lemmas 34, 36, 38, and 39. The rest of the proof follows from the values set for α, β, τ . ◀

References

- 1 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. *Operations Research*, 65(4):1043–1061, 2017. doi:10.1287/opre.2017.1603.
- 2 Julius Borcea, Petter Branden, and Thomas M. Liggett. Negative dependence and the geometry of polynomials. *Journal of American Mathematical Society*, 22:521–567, 2009.
- 3 S. C. Boyd and William R. Pulleyblank. Optimizing over the subtour polytope of the travelling salesman problem. *Math. Program.*, 49:163–187, 1991. doi:10.1007/BF01588786.
- 4 Petter Brändén, Jonathan Leake, and Igor Pak. Lower bounds for contingency tables via lorentzian polynomials. *Israel Journal of Mathematics*, 253(1):43–90, March 2023. doi:10.1007/s11856-022-2364-9.
- 5 Nicos Christofides. Worst case analysis of a new heuristic for the traveling salesman problem. Report 388, Graduate School of Industrial Administration, Carnegie-Mellon University, Pittsburgh, PA, 1976.
- 6 G Dantzig, R Fulkerson, and S Johnson. Solution of a large-scale traveling-salesman problem. *Operations Research*, 2:393–410, 1954. doi:10.1287/OPRE.2.4.393.

- 7 E.A. Dinits, A.V. Karzanov, and M.V. Lomonosov. On the structure of a family of minimal weighted cuts in graphs. *Studies in Discrete Mathematics (in Russian)*, ed. A.A. Fridman, 290–306, Nauka (Moskva), 1976.
- 8 Jack Edmonds and Ellis L. Johnson. Matching, euler tours and the chinese postman. *Mathematical Programming*, 5(1):88–124, 1973. doi:10.1007/BF01580113.
- 9 Tomás Feder and Milena Mihail. Balanced matroids. In *Proceedings of the twenty-fourth annual ACM symposium on Theory of Computing*, pages 26–38, New York, NY, USA, 1992. ACM. doi:10.1145/129712.129716.
- 10 Tamás Fleiner and András Frank. A quick proof for the cactus representation of mincuts. Technical Report QP-2009-03, Egerváry Research Group, Budapest, 2009. www.cs.elte.hu/egres.
- 11 Michel X. Goemans and Dimitris Bertsimas. Survivable networks, linear programming relaxations and the parsimonious property. *Math. Program.*, 60:145–166, 1993. doi:10.1007/BF01580607.
- 12 Anupam Gupta, Euiwoong Lee, Jason Li, Marcin Mucha, Heather Newman, and Sherry Sarkar. Matroid-based TSP rounding for half-integral solutions. In Karen Aardal and Laura Sanità, editors, *Integer Programming and Combinatorial Optimization*, number 13265 in Lecture Notes in Computer Science, pages 305–318, 2022. See also [arXiv:2111.09290](https://arxiv.org/abs/2111.09290).
- 13 Leonid Gurvits. Hyperbolic polynomials approach to van der waerden/schrijver-valiant like conjectures: sharper bounds, simpler proofs and algorithmic applications. In Jon M. Kleinberg, editor, *STOC*, pages 417–426. ACM, 2006. doi:10.1145/1132516.1132578.
- 14 Leonid Gurvits. Van der waerden/schrijver-valiant like conjectures and stable (aka hyperbolic) homogeneous polynomials: One theorem for all. *Electr. J. Comb.*, 15(1), 2008. URL: http://www.combinatorics.org/Volume_15/Abstracts/v15i1r66.html.
- 15 Leonid Gurvits. Boolean matrices with prescribed row/column sums and stable homogeneous polynomials: Combinatorial and algorithmic applications. *Information and Computation*, 240:42–55, 2015. MFCS 2013. doi:10.1016/j.ic.2014.09.007.
- 16 Leonid Gurvits, Nathan Klein, and Jonathan Leake. From Trees to Polynomials and Back Again: New Capacity Bounds with Applications to TSP. In Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson, editors, *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, volume 297 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 79:1–79:20, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICALP.2024.79.
- 17 Leonid Gurvits and Jonathan Leake. Capacity lower bounds via productization. In *STOC*, pages 847–858, 2021. doi:10.1145/3406325.3451105.
- 18 Arash Haddadan and Alantha Newman. Towards improving christofides algorithm for half-integer TSP. In Michael A. Bender, Ola Svensson, and Grzegorz Herman, editors, *ESA*, volume 144 of *LIPIcs*, pages 56:1–56:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ESA.2019.56.
- 19 M. Held and R.M. Karp. The traveling salesman problem and minimum spanning trees. *Operations Research*, 18:1138–1162, 1970. doi:10.1287/OPRE.18.6.1138.
- 20 W. Hoeffding. On the distribution of the number of successes in independent trials. *Ann. Math. Statist.*, 27:713–721, 1956.
- 21 Billy Jin, Nathan Klein, and David P. Williamson. A lower bound for the max entropy algorithm for tsp. In Jens Vygen and Jarosław Byrka, editors, *Integer Programming and Combinatorial Optimization*, pages 238–251, Cham, 2024. Springer Nature Switzerland. doi:10.1007/978-3-031-59835-7_18.
- 22 Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. An improved approximation algorithm for TSP in the half integral case. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *STOC*, pages 28–39. ACM, 2020. doi:10.1145/3357713.3384273.
- 23 Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. A (slightly) improved approximation algorithm for metric tsp. In *STOC*. ACM, 2021. doi:10.1145/3406325.3451009.

- 24 Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. A deterministic better-than- $3/2$ approximation algorithm for metric tsp. In *Integer Programming and Combinatorial Optimization: 24th International Conference, IPCO 2023, Madison, WI, USA, June 21–23, 2023, Proceedings*, pages 261–274. Springer-Verlag, 2023. doi:10.1007/978-3-031-32726-1_19.
- 25 Marek Karpinski, Michael Lampis, and Richard Schmied. New inapproximability bounds for tsp. *Journal of Computer and System Sciences*, 81(8):1665–1677, 2015. doi:10.1016/J.JCSS.2015.06.003.
- 26 Nathan Klein. *Finding Structure in Entropy: Improved Approximation Algorithms for TSP and other Graph Problems*. PhD thesis, University of Washington, USA, 2024. URL: <http://hdl.handle.net/1773/51140>.
- 27 Tobias Mömke and Ola Svensson. Removing and adding edges for the traveling salesman problem. *Journal of the ACM*, 63, 2016. Article 2. doi:10.1145/2739008.
- 28 M Mucha. $\frac{13}{9}$ -approximation for graphic tsp. In *STACS*, pages 30–41, 2012.
- 29 Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. A randomized rounding approach to the traveling salesman problem. In *FOCS*, pages 550–559. IEEE Computer Society, 2011. doi:10.1109/FOCS.2011.80.
- 30 Jim Pitman. Probabilistic bounds on the coefficients of polynomials with only real zeros. *J. Comb. Theory Ser. A*, 77:279–303, February 1997. doi:10.1006/JCTA.1997.2747.
- 31 Frans Schalekamp, David P. Williamson, and Anke van Zuylen. 2-matchings, the traveling salesman problem, and the subtour lp: A proof of the boyd-carr conjecture. *Mathematics of Operations Research*, 39(2):403–417, 2013. doi:10.1287/MOOR.2013.0608.
- 32 András Sebő and Jens Vygen. Shorter tours by nicer ears:. CoRR abs/1201.1870, 2012.
- 33 A. I. Serdyukov. O nekotorykh ekstremal’nykh obkhodakh v grafakh. *Upravlyaemye sistemy*, 17:76–79, 1978. URL: http://nas1.math.nsc.ru/aim/journals/us/us17/us17_007.pdf.
- 34 Damian Straszak and Nisheeth K. Vishnoi. Maximum entropy distributions: Bit complexity and stability. In Alina Beygelzimer and Daniel Hsu, editors, *COLT*, volume 99 of *Proceedings of Machine Learning Research*, pages 2861–2891. PMLR, 2019. URL: <http://proceedings.mlr.press/v99/straszak19a.html>.
- 35 David P. Williamson. Analysis of the Held–Karp heuristic for the traveling salesman problem. Master’s thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, June 1990. URL: <https://dspace.mit.edu/handle/1721.1/149691>.
- 36 Rico Zenklusen. A 1.5-approximation for path TSP. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’19, pages 1539–1549, USA, 2019. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611975482.93.