Matroid-Based TSP Rounding for Half-Integral Solutions

Anupam Gupta[1], Euiwoong Lee[3], Jason Li[2], Marcin Mucha[4], Heather Newman[1](✉), and Sherry Sarkar[1]

[1] Carnegie Mellon University, Pittsburgh, USA
hanewman@andrew.cmu.edu
[2] Simons Institute for the Theory of Computing and UC Berkeley, Berkeley, USA
[3] University of Michigan, Ann Arbor, USA
[4] University of Warsaw, Warsaw, Poland

Abstract. We show how to round any half-integral solution to the subtour-elimination relaxation for the TSP, while losing a less-than-1.5 factor. Such a rounding algorithm was recently given by Karlin, Klein, and Oveis Gharan based on sampling from max-entropy distributions. We build on an approach of Haddadan and Newman to show how sampling from the matroid intersection polytope, and a new use of max-entropy sampling, can give better guarantees.

1 Introduction

The (symmetric) traveling salesman problem asks: given an graph $G = (V, E)$ with edge-lengths $c_e \geq 0$, find the shortest tour that visits all vertices at least once. The Christofides-Serdyukov algorithm [1,10] gives a $3/2$-approximation to this APX-hard problem; this was recently improved to a $(3/2 - \varepsilon)$-approximation by the breakthrough work of Karlin, Klein, and Oveis Gharan, where $\varepsilon > 0$ [7]. A related question is: *what is the integrality gap of the subtour-elimination polytope relaxation for the TSP*? Wolsey had adapted the Christofides-Serdyukov analysis to show an upper bound of $3/2$ [12] (also [11]), and there exists a lower bound of $4/3$. Building on their above-mentioned work, Karlin, Klein, and Oveis Gharan gave an integrality gap of $1.5 - \varepsilon'$ for another small constant $\varepsilon' > 0$ [5], thereby making the first progress towards the conjectured optimal value of $4/3$ in nearly half a century.

Both these recent results are based on a randomized version of the Christofides-Serdyukov algorithm proposed by Oveis Gharan, Saberi, and Singh [8]. This algorithm first samples a spanning tree (plus perhaps one edge) from the *max-entropy distribution* with marginals matching the LP solution, and adds an O-join on the odd-degree vertices O in it, thereby getting an Eulerian spanning subgraph. Since the first step has expected cost equal to that of the LP solution, these works then bound the cost of this O-join by strictly less than half the optimal value, or the LP value. The proof uses a cactus-like decomposition of the min-cuts of the graph with respect to the values x_e, like in [8].

K. Aardal and L. Sanitá (Eds.): IPCO 2022, LNCS 13265, pp. 305–318, 2022.
https://doi.org/10.1007/978-3-031-06901-7_23

Given the $3/2$ barrier has been broken, we can ask: what other techniques can be effective here? How can we make further progress? These questions are interesting even for cases where the LP has additional structure. The half-integral cases (i.e., points for which $x_e \in \{0, 1/2, 1\}$ for all e) are particularly interesting due to the Schalekamp, Williamson, and van Zuylen conjecture, which says that the integrality gap is achieved on instances where the LP has optimal half-integral solutions [9]. The team of Karlin, Klein, and Oveis Gharan first used their max-entropy approach to get an integrality gap of 1.49993 for half-integral LP solutions [6], before they moved on to the general case in [7] and obtained an integrality gap of $1.5 - \varepsilon$; the latter improvement is considerably smaller than in the half-integral case. It is natural to ask: can we do better for half-integral instances?

In this paper, we answer this question affirmatively. We show how to get tours of expected cost of ≈ 1.499 times the linear program value using an algorithm based just on matroid intersection techniques. Moreover, some of these ideas can also strengthen the max-entropy sampling approach in the half-integral case. While the matroid intersection approach and the strengthened max-entropy approach each separately yield improvements over the bound in [6], the improvement obtained by combining these two approaches is slightly better:

Theorem 1.1. *Let x be a half-integral solution to the subtour elimination polytope with cost $c(x)$. There is a randomized algorithm that rounds x to an integral solution whose cost is at most $(1.5 - \varepsilon) \cdot c(x)$, where $\varepsilon = 0.001695$.*

We view our work as showing a proof-of-concept of the efficacy of combinatorial techniques (matroid intersection, and flow-based charging arguments) in getting an improvement for the half-integral case. We hope that these techniques, ideally combined with max-entropy sampling techniques, can give further progress on this central problem.

Our Techniques. The algorithm is again in the Christofides-Serdyukov framework. It is easiest to explain for the case where the graph (a) has an even number of vertices, and (b) has no (non-trivial) proper min-cuts with respect to the LP solution values x_e—specifically, the only sets for which $x(\partial S) = 2$ correspond to the singleton cuts. Here, our goal is that each edge is "even" with some probability: i.e., both of its endpoints have even degree with probability $p > 0$. In this case we use an idea due to Haddadan and Newman [3]: we *shift* and get a $\{1/3, 1\}$-valued solution y to the subtour elimination polytope K_{TSP}. Specifically, we find a random perfect matching M in the support of x, and set $y_e = 1$ for $e \in M$, and $1/3$ otherwise, thereby ensuring $\mathbb{E}[y] = x$. To pick a random tree from this shifted distribution y, we do one of the following:

1. We pick a random "independent" set M' of matching edges (so that no edge in E is incident to two edges of M'). For each $e' \in M'$, we place partition matroid constraints enforcing that exactly one edge is picked at each endpoint—which, along with e' itself, gives degree 2 and thereby makes the edge even as desired. Finding spanning trees subject to another matroid constraint can be implemented using matroid intersection.
2. Or, instead we sample a random spanning tree from the max-entropy distribution, with marginals being the shifted value y. (In contrast, [6] sample trees

from x itself; our shifting allows us to get stronger notions of evenness than they do: e.g., we can show that every edge is "even-at-last" with constant probability, as opposed to having at least one even-at-last edge in each tight cut with some probabiltiy.)

(Our algorithm randomizes between the two samplers to achieve the best guarantees.) For the O-join step, it suffices to give fractional values z_e to edges so that for every odd cut in T, the z-mass leaving the cut is at least 1. In the special case we consider, each edge only participates in two min-cuts—those corresponding to its two endpoints. So set $z_e = x_e/3$ if e is even, and $x_e/2$ if not; the only cuts with $z(\partial S) < 1$ are minimum cuts, and these cuts will not show up as O-join constraints, due to evenness. For this setting, if an edge is even with probability p, we get a $(3/2 - p/6)$-approximation!

It remains to get rid of the two simplifying assumptions. To sample trees when $|V|$ is odd (an open question from [3]), we add a new vertex to fix the parity, and perform local surgery on the solution to get a new TSP solution and reduce to the even case. The challenge here is to show that the losses incurred are small, and hence each edge is still even with constant probability.

Finally, what if there are proper tight sets S, i.e., where $x(\partial S) = 2$? We use the cactus decomposition of a graph (also used in [6,8]) to sample spanning trees from pieces of G with no proper min-cuts, and stitch these trees together. These pieces are formed by contracting sets of vertices in G, and have a hierachical structure. Moreover, each such piece is either of the form above (a graph with no proper min-cuts) for which we have already seen samplers, or else it is a double-edged cycle (which is easily sampled from). Since each edge may now lie in many min-cuts, we no longer just want an edge to have both endpoints be even. Instead, we use an idea from [6] that uses the hierarchical structure on the pieces considered above. Every edge of the graph is "settled" at exactly one of these pieces, and we ask for both of its endpoints to have even degree *in the piece at which it is settled*. The z_e value of such an edge may be lowered from an initial value of $x_e/2$ in the O-join without affecting constraints corresponding to cuts *in the piece at which it is settled*.

Since cuts at other levels of the hierarchy may now be deficient because of the lower values of z_e, we may need to increase the z_f values for other "lower" edges f to satisfy these deficient cuts. This last part requires a charging argument, showing that each edge e has z_e that is strictly smaller than $x_e/2$ in expectation. For our samplers, the naïve approach of distributing charge uniformly as in [6] does not work, so we instead formulate this charging as a flow problem.

Due to lack of space we present the simpler samplers and the main algorithm here, and defer many of the proofs and the details to the full version.

2 Notation and Preliminaries

Given a multigraph $G = (V, E)$, and a set $S \subseteq V$, let ∂S denote the cut consisting of the edges connecting S to $V \setminus S$; S and $\bar{S} := V \setminus S$ are called *shores* of the cut. A subset $S \subseteq V$ is *proper* if $1 < |S| < |V| - 1$; a cut ∂S is called *proper* if the set S is a proper subset. A set S is *tight* if $|\partial S|$ equals the size of the minimum edge-cut in G. Two sets S and S' are *crossing* if $S \cap S'$, $S \setminus S'$, $S' \setminus S$, and $V \setminus (S \cup S')$ are all non-empty.

Define the *subtour elimination polytope* $K_{TSP}(G) \subseteq \mathbb{R}^{|E|}$:

$$\{x \geq 0 \mid x(\partial v) = 2 \; \forall v \in V, x(\partial S) \geq 2 \; \forall \text{ proper } S\}. \tag{1}$$

Let x be half-integral and feasible for (1). W.l.o.g. we can focus on solutions with $x_e = 1/2$ for each $e \in E$, doubling edges if necessary. The support graph G is then a 4-regular 4-edge-connected (henceforth 4EC) multigraph.

Fact 2.1. *If $x \in K_{TSP}(G)$, then $x|_{E(V(G)\setminus\{r\})}$ is in the spanning tree polytope $K_{spT}(G[V(G) \setminus \{r\}])$ for any $r \in V$, and $x/2$ is in the perfect matching polytope $K_{PM}(G)$ (when $|V(G)|$ is even) and in the O-join dominator polytope $K_{join}(G, O)$, $O \subseteq V(G)$, $|O|$ even, given by:*

$$\{z \geq 0 \mid z(\partial S) \geq 1 \; \forall S \subseteq V, |S \cap O| \text{ odd.}\}.$$

Lemma 2.1. *Consider a sub-partition $\mathcal{P} = \{P_1, P_2, \ldots, P_t\}$ of the edge set of G. Let x be a fractional solution to the spanning tree polytope that satisfies $x(P_i) \leq 1$ for all $i \in [t]$. The integrality of the matroid intersection polytope implies that we can efficiently sample from a probability distribution $\mathcal{D}$ over spanning trees which contain at most one edge from each of the parts P_i, such that $\mathbb{P}_{T \leftarrow \mathcal{D}}[e \in T] = x_e$.*

Let z be in the relative interior of the spanning tree polytope. The *max-entropy distribution* is a distribution μ of spanning trees that maximizes the entropy of μ subject to $\mathbb{P}(e \in T) = z_e$ [8]. It is a λ-uniform spanning tree distribution and thus is *strongly Rayleigh (SR).*

Theorem 2.2 (Negative Correlation [8]). *Let μ be an SR distribution on spanning trees.*

1. *Let S be a set of edges and $X_S = |S \cap T|$, where $T \sim \mu$. Then, $X_S \sim \sum_{i=1}^{|S|} Y_i$, where the Y_i are independent Bernoulli random variables with success probabilities p_i and $\sum_i p_i = \mathbb{E}[X_S]$.*
2. *For any set of edges S and $e \notin S$,*
   (i) $\mathbb{E}_\mu[X_S] \leq \mathbb{E}_\mu[X_S \mid X_e = 0] \leq \mathbb{E}_\mu[X_S] + \mathbb{P}_\mu(e \in T)$, and
   (ii) $\mathbb{E}_\mu[X_S] - 1 + \mathbb{P}_\mu(e \in T) \leq \mathbb{E}_\mu[X_S \mid X_e = 1] \leq \mathbb{E}_\mu[X_S]$.

Theorem 2.3 ([4], Corollary 2.1). *Let $g : \{1, \ldots, m\} \to \mathbb{R}$ and $0 \leq p \leq m$. Let $B_1, \ldots, B_m$ be Bernoulli r.v.s with probabilities $p_1^*, \ldots, p_m^*$ that maximize (or minimize) $\mathbb{E}[g(B_1 + \ldots + B_m)]$ over all possible success probabilities p_i for B_i for which $p_1 + \cdots + p_m = p$. Then $\{p_1^*, \ldots, p_m^*\} \in \{0, x, 1\}$ for some $x \in (0, 1)$.*

3 Samplers

We now describe the MaxEnt and MatInt samplers for graphs that contain no proper min-cuts, and give bounds on certain correlations between edges that will be used in Sect. 5 to prove that every edge is "even" with constant probability.

For lack of space, we focus on the case where $|V|$ is even; the case of $|V|$ being odd is slightly more technical; please see the full version of the paper.

Suppose the graph $H = (V, E)$ is 4-regular and 4EC, contains at least four vertices, and has no proper min-cuts. H is a simple graph, because parallel edges between u, v would mean that $\partial(\{u, v\})$ is a proper min-cut. Also, all proper cuts have six or more edges. We are given a dedicated *external* vertex $r \in V(H)$; the vertices $I := V \setminus \{r\}$ are called *internal*. (In future sections, r will be given by a cut hierarchy.) Call the edges in ∂r *external edges*; all other edges are *internal*. An internal vertex is called a *boundary vertex* if it is adjacent to r. An edge is said to be *special* if both of its endpoints are non-boundary vertices.

We show two ways to sample a spanning tree on $H[I]$, the graph induced on the internal vertices, being faithful to the marginals x_e, i.e., $\mathbb{P}_T(e \in T) = x_e$ for all $e \in E(H) \setminus \partial r$. Moreover, we want that for each internal edge, both its endpoints have even degree in T with constant probability. This property will allow us to lower the cost of the O-join in Sect. 6. While both samplers will satisfy this property, each will do better in certain cases. The MatInt sampler targets special edges; it allows us to randomly "hand-pick" edges of this form and enforce that both of its endpoints have degree 2 in the tree. The MaxEnt sampler, on the other hand, relies on maximizing the randomness of the spanning tree sampled (subject to being faithful to the marginals); negative correlation properties allow us to obtain the evenness property, specifically, better probabilities than MatInt for non-special edges, and a worse one for the special edges.

Our samplers will depend on the parity of $|V|$: when $|V|$ is even, the MatInt sampler in Sect. 3.1 was given by [3, Theorem 13]. They left the case of odd $|V|$ as an open problem, which we solve.

3.1 Samplers for Even $|V(H)|$

Since H is 4-regular and 4EC and $|V(H)|$ is even, setting a value of $1/4 = x_e/2$ on each edge gives a solution to $K_{PM}(H)$ by Fact 2.1.

1. Sample a perfect matching M s.t. $\mathbb{P}(e \in M) = 1/4 = x_e/2$ for all $e \in E(H)$.
2. Define a new fractional solution y (that depends on M): set $y_e = 1$ for $e \in M$, and $y_e = 1/3$ otherwise. We have $y \in K_{TSP}(H)$ (and hence $y|_I \in K_{spT}(H[I])$ by Fact 2.1): indeed, each vertex has $y(\partial v) = 1 + 3 \cdot 1/3 = 2$ because M is a perfect matching. Moreover, every proper cut U in H has at least six edges, so $y(\partial U) \geq |\partial U| \cdot 1/3 \geq 2$. Furthermore,

$$\mathbb{E}_M[y_e] = 1/4 \cdot 1 + 3/4 \cdot 1/3 = 1/2 = x_e. \tag{2}$$

3. Sample a spanning tree faithful to the marginals y, using one of two samplers:
   (a) MaxEnt Sampler: Sample from the max-entropy distribution on spanning trees with marginals y. (Since y may not be in the relative interior of the spanning tree polytope, contract the 1-valued edges to obtain

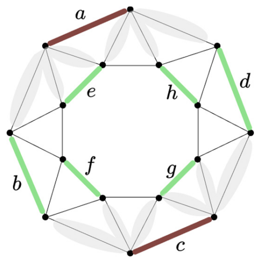

Fig. 1. The matching M consists of the green and brown edges; one possible choice of M' has edges $\{a, c\}$ in brown, and the constraints are placed on the edges adjacent to those in M' (marked in gray).

a 6-regular, $1/3$-uniform solution. This may have nontrivial min-cuts, so once again use a cactus hierarchy to decompose the graph into pieces (see Sect. 4.1); the induced solution on each piece is in the relative interior of the spanning tree polytope. For each piece, sample a λ-uniform spanning tree that preserves marginals, and then stitch these trees together.)

(b) MatInt Sampler:
  i. Color the edges of M using 7 colors such that no edge of H is adjacent to two edges of M having the same color; e.g., greedily 7-color the 6-regular graph H/M. Let M' be one of these color classes picked uniformly at random. So $\mathbb{P}(e \in M') = 1/28$, and $\mathbb{P}(\partial v \cap M' \neq \varnothing) = 1/7$ see Fig. 1.
  ii. For each edge $e = uv \in M'$, let L_{uv} and R_{uv} be the sets of edges incident at u and v other than e. Note that $|L_{uv}| = |R_{uv}| = 3$. Place partition matroid constraints $y(L_{uv}) \leq 1$ and $y(R_{uv}) \leq 1$ on each of these sets. Finally, restrict the partition constraints to the internal edges of H; this means some of these constraints are no longer tight for the solution y.

(c) Given the sub-matching $M' \subseteq M$, and the partition matroid $\mathcal{M}$ on the internal edges defined using M', use Lemma 2.1 to sample a tree on $H[V \setminus \{r\}]$ (i.e., on the internal vertices and edges of H) with marginals y_e, subject to this partition matroid $\mathcal{M}$.

Conditioned on the matching M, we have $\mathbb{P}(e \in T \mid M) = y_e$; now using (2), we have $\mathbb{P}(e \in T) = x_e$ for all $e \in (E \setminus \partial r)$.

The main idea for the odd case is to duplicate the external vertex with a pair of parallel edges between these copies. Since this gives a graph with proper min-cuts, we cannot apply shifting naively. Instead we perform "local surgery" on the LP solution to get a feasible fractional spanning tree. Showing that these changes still give us a tree with good evenness properties requires some care, and the ideas are deferred to the full version for lack of space.

3.2 Correlation Properties of Samplers

Let T be a tree sampled using either the MatInt or the MaxEnt sampler. The following claims will be used to prove the evenness property in Sect. 5. Each table gives lower bounds on the corresponding probabilities for each sampler.

Lemma 3.1. *If f, g are internal edges incident to a vertex v, then*

Probability Statement	MatInt	MaxEnt
$\mathbb{P}(\lvert T \cap \{f, g\}\rvert = 2)$	$1/9$	$1/9$
$\mathbb{P}(T \cap \{f, g\} = \{f\})$	$1/9$	$12/72$

Lemma 3.2. *If edges e, f, g, h incident to a vertex v are all internal, then*

Probability Statement	MatInt	MaxEnt
$\mathbb{P}(\lvert T \cap \{e, f, g, h\}\rvert = 2)$	$2/21$	$8/27$

Lemma 3.3. *For an internal edge $e = uv$:*
(a) if both endpoints are non-boundary vertices, then

Probability Statement	MatInt	MaxEnt
$\mathbb{P}(\lvert\partial_T(u)\rvert = \lvert\partial_T(v)\rvert = 2)$	$1/36$	$128/6561$

(b) if both u, v are boundary vertices, then

Probability Statement	MatInt	MaxEnt
$\mathbb{P}$(exactly one of u, v has odd degree in T)	$1/9$	$5/18$

To give a sense of the techniques, we give the proof for the last statement above when $|V(H)|$ is even. The other proofs are similar in flavor, please see the full version.

Proof (Lemma 3.3a, Even Case). The MatInt *claims:* The event happens when $e \in M'$, which happens w.p. $1/28$, which is at least $1/36$.
The MaxEnt *claims:* Condition on $e \in M$. Let $S_1 = \partial(u) \setminus e$ and $S_2 = \partial(v) \setminus e$. Denote $S_1 = \{a, b, c\}$. Lower bound $\mathbb{P}(|S_1 \cap T| = 1)$ using Theorem 2.3: $\mathbb{E}[|S_1 \cap T|] = 3 \cdot 1/3 = 1$, so $\mathbb{P}(|S_1 \cap T| = 1) \geq 3 \cdot 1/3 \cdot (2/3)^{2/3} = 4/9$. Consider the distribution over the edges in S_2 conditioned on $a \in T$; this distribution is also SR. By Theorem 2.2, $1/3 \leq \mathbb{E}[X_{S_2} \mid X_a = 1] \leq 1$. Applying Theorem 2.2 twice more,

$$1/3 \leq \mathbb{E}[X_{S_2} \mid X_a = 1, X_{b,c} = 0] \leq 1 + 1/3 + 1/3 = 5/3.$$

By Theorem 2.3, $\mathbb{P}(X_{S_2} = 1 \mid X_a = 1, X_{b,c} = 0) \geq 3 \cdot 1/9 \cdot (8/9)^2 = 64/243$. Using symmetry, we obtain $\mathbb{P}(X_{S_2} = 1 \wedge X_{S_1} = 1 \wedge e \in M) \geq 64/243 \cdot 4/9 \cdot 1/4 \geq 128/6561$.

4 Sampling Algorithm for General Solutions

Now that we can sample a spanning tree from a graph with no proper min-cuts, we introduce the algorithm to sample a spanning tree plus one edge (an r_0 tree) from a 4-regular, 4EC graph, perhaps with proper min-cuts.

W.l.o.g., assume that the graph $G = (V, E)$ has a set of three special vertices $\{r_0, u_0, v_0\}$, with each pair r_0, u_0 and r_0, v_0 having a pair of edges between them (used in line 18).(We can introduce dummy nodes to ensure this property, which is for simplicity—it guarantees that the top set in the cut hierarchy is a cycle set.) Define a *double cycle* to be a cycle graph in which each edge is replaced by a pair of parallel edges, and call each such pair *partner edges*.

Algorithm 1. Sampling Algorithm for a Half-Integral Solution

1: **let** G be the support graph of a half-integral solution x.
2: **let** $T = \varnothing$.
3: **while** $\exists$ a proper tight set of G not crossed by another proper tight set **do**
4:     **let** S be a minimal such set (and choose S such that $r_0 \notin S$).
5:     Define $G' = G/(V \setminus S)$.
6:     **if** G' is a double cycle **then**
7:         Label S a *cycle set*.
8:         **sample** a random edge from each set of partner edges in $G[S]$; add these edges to T.
9:     **else** // *G' has no proper min-cuts (Lemma 4.1).*
10:         Label S a *degree set*.
11:         **if** $G' = K_5$ **then**
12:             **sample** a random path on $G[S]$
13:         **else**
14:             W.p. λ, let μ be the MaxEnt distribution over $E(S)$
15:             W.p. $1 - \lambda$, let μ be the MatInt distribution.
16:             **sample** a spanning tree on $G[S]$ from μ and add its edges to T.
17:     **let** $G = G/S$
18: Due to r_0, u_0, v_0, at this point G is a double cycle (Lemma 4.1). Sample one edge between each pair of adjacent vertices in G.

As in [6], we refer to the sets in line 4 as *critical sets*. The algorithm samples from the same pieces as in [6], with the key differences being randomizing between the MatInt and MaxEnt samplers as well as a critical optimization for K_5's (the latter will become clear later and in the full version of the paper).

Lemma 4.1. *Algorithm 1 is well-defined: In every iteration of Algorithm 1, G' is either a double cycle or a graph with no proper min-cuts, and graph remaining at the end of the algorithm (line 18) is a double cycle.*

We will prove the following theorem in Sect. 6. This in turn gives Theorem 1.1.

Theorem 4.1. *Let T be the r_0-tree chosen from Algorithm 1, and O be the set of odd degree vertices in T. The expected cost of the minimum cost O-join for T is at most $(1/2 - \varepsilon) \cdot c(x)$.*

4.1 The Cut Hierarchy

To lower the cost of the O-join, we need a complete description of the min-cuts of G, which will be achieved by the implicit hierarchy of sets Algorithm 1 induces. This hierarchical decomposition is the same as the one used in [6]; however, here we give an explicit way to construct the hierarchy of tight cuts. The hierarchy is given by a rooted tree $\mathcal{T} = (V_\mathcal{T}, E_\mathcal{T})$.[1] The node set $V_\mathcal{T}$ corresponds to all critical sets found by the algorithm, along with a root node and leaf nodes labelled the vertices in $V_G \setminus \{r_0\}$. If S is a critical set, we label the node in $V_\mathcal{T}$ with S, where we view $S \subseteq V_G$ and not $V_{G'}$. The root node is labelled $V_G \setminus \{r_0\}$. A node S is a child of S' if $S \subset S'$ and S' is the first superset of S contracted after S in the algorithm. Also, the root node is a parent of all nodes corresponding to critical sets that are not strictly contained in any other critical set. Each leaf node is a child of the smallest critical set that contains it. Observe that vertex sets labelling the children of a node are a partition of the vertex set labelling that node. A node in $V_\mathcal{T}$ is a *cycle* or *degree* node if the corresponding critical set labelling it is a cycle or degree set. (We take the root node as a cycle node. The leaf nodes are not labelled as degree or cycle nodes see Fig. 2.)

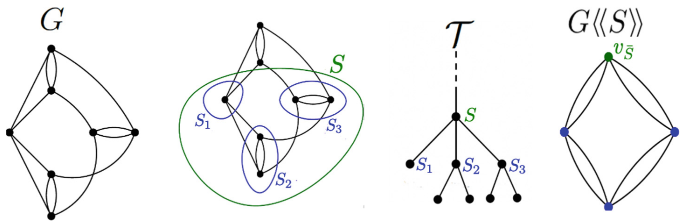

Fig. 2. A portion of the cut hierarchy $\mathcal{T}$ and the local multigraph $G\langle\langle S\rangle\rangle$.

Let $S \subseteq V_G$ be a set labelling a node in $\mathcal{T}$. Define the *local multigraph* $G\langle\langle S\rangle\rangle$ to be the following graph: take G and contract the subsets of V_G labelling the children of S in $\mathcal{T}$ down to single vertices and contract $\bar{S}$ to a single vertex $v_{\bar{S}}$. Remove any self-loops. The vertex $v_{\bar{S}}$ is called the *external vertex*; all other vertices are called *internal vertices.* An internal vertex is called a *boundary vertex* if it is adjacent to the external vertex. The edges in $G\langle\langle S\rangle\rangle \setminus v_{\bar{S}}$ are called *internal edges.* Observe $G\langle\langle S\rangle\rangle$ is precisely the graph G' in line 5 of Algorithm 1 when S is a critical set, and is a double cycle when $S = V_G \setminus \{r_0\}$.

[1] Since there are several graphs under consideration, the vertex set of G is called V_G. Moreover, for clarity, we refer to elements of V_G as vertices, and elements of $V_\mathcal{T}$ as nodes.

Properties of $\mathcal{T}$

1. Let $S \subseteq V_G$ be a set labelling a node in $\mathcal{T}$. If S is a degree node in $\mathcal{T}$, then $G\langle\langle S\rangle\rangle$ has at least five vertices and no proper min-cuts. If S is a cycle node in $\mathcal{T}$, then $G\langle\langle S\rangle\rangle$ is exactly a double cycle. These are by Lemma 4.1.
2. Algorithm 1 can be restated: For each non-leaf node S in $\mathcal{T}$, sample a random path on $G\langle\langle S\rangle\rangle \setminus v_{\bar{S}}$ if it is a double cycle or K_5; otherwise use the MaxEnt or MatInt samplers w.p. λ and $1-\lambda$, respectively, on $G\langle\langle S\rangle\rangle \setminus v_{\bar{S}}$.
3. For a degree set S, the graph $G\langle\langle S\rangle\rangle$ having no proper min-cuts implies that it has no parallel edges. In particular, no vertex has parallel edges to the external vertex in $G\langle\langle S\rangle\rangle$. Hence we get the following:

Corollary 4.1. *For a set S labeling a non-leaf node in $\mathcal{T}$ and any internal vertex $v \in G\langle\langle S\rangle\rangle$: if S is a cycle set then $|\partial v \cap \partial S| \in \{0, 2\}$, and if S is a degree set then $|\partial v \cap \partial S| \in \{0, 1\}$.*

The cactus representation of min-cuts [2] translates to the following complete characterization of the min-cuts of G in terms of local multigraphs.

Lemma 4.2. *Any min-cut in G is either (a) ∂S for some node S in $\mathcal{T}$, or (b) ∂X where X is obtained as follows: for some cycle set S in $\mathcal{T}$, X is the union of vertices corresponding to some contiguous segment of the cycle $G\langle\langle S\rangle\rangle$.*

5 Analysis Part I: The Even-at-Last Property

We now define a notion of evenness for every edge in G that will allow us to reduce the cost of the O-join in Sect. 6. In the case where G has no proper min-cuts, we called an edge *even* if both of its endpoints were even in T. The general definition of evenness extends this idea, but now depends on where an edge belongs in the hierarchy $\mathcal{T}$. Specifically, we say an edge $e \in E(G)$ is *settled* at S if S is the (unique) set such that e is an internal edge of $G\langle\langle S\rangle\rangle$; call S the *last set* of e. If S is a degree or cycle set, we call e a *degree edge* or *cycle edge*.

Let S be the last set of e, and $T\langle\langle S\rangle\rangle$ be the restriction of T to $G\langle\langle S\rangle\rangle$.

1. A degree edge e is called *even-at-last (EAL)* if both its endpoints have even degree in $T\langle\langle S\rangle\rangle$.
2. For a cycle edge $e = uv$, the graph $G\langle\langle S\rangle\rangle \setminus \{v_{\bar{S}}\}$ is a chain of vertices $v_\ell, \ldots, u, v, \ldots, v_r$, with consecutive vertices connected by two parallel edges. Let $C := \{v_\ell, \ldots, u\}$, and $C' := \{v, \ldots, v_r\}$ be a partition of this chain. The cuts ∂C and $\partial C'$ are called the *canonical cuts* for e. Cycle edge e is called *even-at-last (EAL)* if both canonical cuts are crossed an even number of times by $T\langle\langle S\rangle\rangle$; in other words, if there is exactly one edge in $T\langle\langle S\rangle\rangle$ from each of the two pairs of external partner edges leaving v_ℓ and v_r.

Informally, a degree edge is EAL in the general case if it is even in the tree at the level at which it is settled. Let e be settled at a *degree* set S. We say that e is *special* if both of its endpoints are non-boundary vertices in $G\langle\langle S\rangle\rangle$ and *half-special* if exactly one of its endpoints is a boundary vertex in $G\langle\langle S\rangle\rangle$. The key property used in Sect. 6 to reduce each z_e in the fractional O-join is:

Theorem 5.1 (The Even-at-Last Property). *The table below gives lower bounds on the probability that special, half-special, and all other types of degree edges are EAL in each of the two samplers.*

	special	*half-special*	*other degree edges*
MatInt	$1/36$	$1/21$	$1/18$
MaxEnt	$128/6561$	$4/27$	$12/144$

Moreover, a cycle edge is EAL w.p. at least $\lambda \cdot 12/144 + (1-\lambda)\cdot 1/18$.

Proof. Let e be settled at S. Let T_S be the spanning tree sampled on the internal vertices of $G\langle\langle S\rangle\rangle$ (in Algorithm 1, the spanning tree sampled on $G[S]$). We show the proof when S is a degree set.

1. If e is special, then Lemma 3.3(a) gives the bounds in the table.
2. Suppose one of the endpoints of $e = uv$ (say u) is a boundary vertex in S, with edge f incident to u leaving S. By Lemma 3.2, the other endpoint v is even in T_S w.p. $2/21$ for the MatInt sampler and $8/27$ for the MaxEnt sampler. Moreover, the edge f is chosen at a higher level than S and is therefore independent of T_S, and hence can make the degree of u even w.p. $1/2$. Thus e is EAL w.p. $1/21$ for the MatInt sampler and $4/27$ for the MaxEnt sampler.
3. Suppose both endpoints of e are boundary vertices of S, with edges f, g leaving S. Let $q_=$ be the probability that the degrees of vertices u, v in the tree T_S chosen within S have the same parity, and $q_{\neq} = 1 - q_=$. Now, when S is contracted and we choose a r_0-tree T' on the graph G/S consistent with the marginals, let $p_=$ be the probability that either both or neither of f, g are chosen in T', and $p_{\neq} = 1 - p_=$. Hence

$$\mathbb{P}(e\ EAL) = q_{oo}p_{11} + q_{oe}p_{10} + q_{eo}p_{01} + q_{ee}p_{00} = 1/2(p_= q_= + p_{\neq} q_{\neq}), \quad (3)$$

   where q_{oo}, q_{oe}, q_{eo} and q_{ee} correspond to different parity combinations of u and v in T_S and $p_{00}, p_{01}, p_{10}, p_{11}$ correspond to whether f and g are chosen in T'. The second equality follows from symmetry.
   (a) If f, g are settled at different levels, then they are independent. This gives $p_= = p_{\neq} = 1/2$, and hence $\mathbb{P}(e\ EAL) = 1/4$ regardless of the sampler.
   (b) If f, g have the same last set which is a degree set, then by Lemma 3.1 $p_{11}, p_{01}, p_{10} \geq 1/9$. By symmetry, $p_{00} \geq 1/9$. So (3) gives $\mathbb{P}(e\ EAL) \geq 1/9$.
   (c) If f, g have the same last set which is a cycle set, consider the case where f, g are partners, in which case $p_{\neq} = 1$. Now (3) implies that $\mathbb{P}(e\ EAL) = q_{\neq}/2$, which by Lemma 3.3(b) is $\geq 1/2 \cdot 1/9 = 1/18$ in the MatInt sampler, and $1/2 \cdot 5/18 = 5/36$ in the MaxEnt sampler. If f, g are not partners, then they are chosen independently, in which case again $p_= = p_{\neq} = 1/2$, and hence $\mathbb{P}(e\ EAL) = 1/4$.

The proof for cycle sets follows similar lines, and is in the full version.

6 Analysis Part II: The O-Join and Charging

To prove Theorem 4.1, we construct an O-join for the random tree T, and bound its expected cost via a charging argument. The structure of here is similar to [6]; however, we use a flow-based approach to perform the charging instead of the naive one, and also use our stronger property that *every* edge is EAL with constant probability (versus the weaker property obtained in [6] that every tight cut contains an EAL edge with constant probability).

Let O denote the (random set of) odd-degree vertices in T. The dominant of the O-join polytope $K_{join}(G, O)$ is given by

$$\{x(\partial S) \geq 1 \mid \forall S \subseteq V, |\partial S \cap T| \text{ odd.}\}.$$

This polytope is integral, so it suffices to show the *existence* of a fractional O-join solution $z \in K_{join}(G, O)$ with low expected cost. (The expectation is taken over O.)

To construct the fractional O-join z, we begin with $z = x/2$. Notice that $z(\partial S) \geq 1$ is a tight constraint when S is a min-cut. For any e that is EAL in T, we first flip a biased coin to know whether to reduce z_e: the purpose of the coin flips is to "flatten" the probability of reducing z_e to the bound given by Theorem 5.1 on the probability that e is EAL. Now the amount of the reduction in z_e depends on whether e is a degree or cycle edge; the amount is later optimized by the solution to a linear program. In the case of the cycle edge, we are able to reduce an edge by the full $1/12$; degree edges cannot be reduced as drastically.

However, these reductions may make z infeasible, and we need to fix that. Indeed, suppose f is EAL and that we reduce z_f. Say f is settled at S. If S is a degree set, then the only min-cuts of $G\langle\langle S\rangle\rangle$ are the degree cuts. So the only min-cuts that the edge f is part of in $G\langle\langle S\rangle\rangle$ are the degree cuts of its endpoints; call them U, V, in $G\langle\langle S\rangle\rangle$ (U and V are vertices in $G\langle\langle S\rangle\rangle$ representing sets U and V in G). But since $|\partial U \cap T|$ and $|\partial V \cap T|$ are both even by definition of EAL, we need not worry that reducing z_f causes $z(\partial U) \geq 1$ and $z(\partial V) \geq 1$ to be violated. Likewise, if S is a cycle set, then by definition of EAL all min-cuts S' in $G\langle\langle S\rangle\rangle$ containing e have $|\partial S' \cap T|$ even, so again we need not worry.

Since f is only an internal edge for its last set S, the only cuts S' for which the constraint $z(\partial S') \geq 1$, $|\partial S' \cap T|$ odd, may be violated as a result of reducing z_f are cuts represented in lower levels of the hierarchy. Specifically, lef f be an external edge for some $G\langle\langle X\rangle\rangle$ (meaning X is lower in the hierarchy $\mathcal{T}$ than S) and S' be a min-cut of $G\langle\langle X\rangle\rangle$ (either a degree cut or a canonical cut). By Lemma 4.2, cuts of the form S' are the only cuts that may be deficient as a result of reducing z_f. Call the internal edges of $\partial S'$ *lower edges*. When z_f is reduced and $|\partial S' \cap T|$ is odd, we must distribute an increase (charge) over the lower edges totalling the amount by which z_f is reduced, so that $z(\partial S') = 1$.

Fix edge e, say with last set S. We need to show that in expectation the charge it receives from external edges in $G\langle\langle S\rangle\rangle$ is strictly less than the initial expected reduction to z_e. External edges bring in charge to internal edges and Corollary 4.1 says that every vertex in a critical set can have at 0, 1, or 2 external edges (and there are 4 external edges). When e is a cycle edge, we distribute charge from an external edge evenly between e and its partner. When e is a degree edge, charge will be optimally distributed according to a maximum-flow solution. Specifically, in order to minimize the maximum charge given on any edge, we bipartize $G\langle\langle S\rangle\rangle$ into $H = (L, R)$: the vertices in L and R represent external and internal edges, respectively, and the edges of H are those pairs of edges in $G\langle\langle S\rangle\rangle$ which are adjacent. Each vertex in L releases a unit of charge and each vertex in R absorbs at most c units of charge.

The problem of optimally distributing charge reduces to finding the smallest constant c such that there exists a flow with capacity at most c. An argument based on Hall's condition characterizes precisely when a flow distributing c units of charge to internal edges exists. In order to optimize the constant c (found to be $1/2$), the case where $G\langle\langle S\rangle\rangle = K_5$ happens to be a bottleneck; hence, we treat the K_5 case separately in order to gain from the max-flow formulation, by choosing a very natural sampling method for it, and then reducing its edges differently from those of other degree sets.

Acknowledgments. AG supported in part by NSF awards CCF-1907820, CCF-1955785, and CCF-2006953. MM partially supported by the ERC CoG grant TUgbOAT no 772346.

References

1. Christofides, N.: Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report 338, Graduate School of Industrial Administration, Carnegie-Mellon University, Pittsburgh (1976)
2. Fleiner, T., Frank, A.: A quick proof for the cactus representation of mincuts. Technical report Quick-proof No. 2009–03, EGRES (2009)
3. Haddadan, A., Newman, A.: Towards improving Christofides algorithm for half-integer TSP. In: ESA, pp. 56:1–56:12 (2019). https://doi.org/10.4230/LIPIcs.ESA.2019.56
4. Hoeffding, W.: On the distribution of the number of successes in independent trials. Ann. Math. Stat. **27**(3), 713–721 (1956)
5. Karlin, A., Klein, N., Oveis Gharan, S.: A (slightly) improved bound on the integrality gap of the subtour LP for TSP. CoRR abs/2105.10043 (2021). https://arxiv.org/abs/2105.10043
6. Karlin, A.R., Klein, N., Oveis Gharan, S.: An improved approximation algorithm for TSP in the half integral case. In: STOC, pp. 28–39. ACM (2020). https://doi.org/10.1145/3357713.3384273
7. Karlin, A.R., Klein, N., Oveis Gharan, S.: A (slightly) improved approximation algorithm for metric TSP. CoRR abs/2007.01409 (2020). https://arxiv.org/abs/2007.01409
8. Oveis Gharan, S., Saberi, A., Singh, M.: A randomized rounding approach to the traveling salesman problem. In: FOCS, pp. 550–559 (2011). https://doi.org/10.1109/FOCS.2011.80
9. Schalekamp, F., Williamson, D.P., van Zuylen, A.: 2-matchings, the traveling salesman problem, and the subtour LP: a proof of the boyd-carr conjecture. Math. Oper. Res. **39**(2), 403–417 (2014). https://doi.org/10.1287/moor.2013.0608

10. Serdyukov, A.: O nekotorykh ekstremal'nykh obkhodakh v grafakh. Upravlyaemye sistemy, **17**, 76–79 (1978). http://nas1.math.nsc.ru/aim/journals/us/us17/us17_007.pdf
11. Shmoys, D.B., Williamson, D.P.: Analyzing the Held-Karp TSP bound: a monotonicity property with application. Info. Proc. Let. **35**(6), 281–285 (1990)
12. Wolsey, L.A.: Heuristic analysis, linear programming and branch and bound. In: Combinatorial Optimization II, Mathematical Programming Studies, vol. 13, pp. 121–134. Springer (1980). https://doi.org/10.1007/BFb0120913