A 4/3-Approximation Algorithm for Half-Integral Cycle Cut Instances of the TSP

Billy Jin[1](✉), Nathan Klein[2], and David P. Williamson[1]

[1] Cornell University, Ithaca, USA
{bzj3,davidpwilliamson}@cornell.edu
[2] University of Washington, Seattle, USA
nwklein@cs.washington.edu

Abstract. A long-standing conjecture for the traveling salesman problem (TSP) states that the integrality gap of the standard linear programming relaxation of the TSP (sometimes called the Subtour LP or the Held-Karp bound) is at most 4/3 for symmetric instances of the TSP obeying the triangle inequality. In this paper we consider the half-integral case, in which a feasible solution to the LP has solution values in $\{0, 1/2, 1\}$. Karlin, Klein, and Oveis Gharan [9], in a breakthrough result, were able to show that in the half-integral case, the integrality gap is at most 1.49993; Gupta et al. [6] showed a slight improvement of this result to 1.4983.

Both of these papers consider a hierarchy of critical tight sets in the support graph of the LP solution, in which some of the sets correspond to *cycle cuts* and the others to *degree cuts*. Here we show that if all the sets in the hierarchy correspond to cycle cuts, then we can find a distribution of tours whose expected cost is at most 4/3 times the value of the half-integral LP solution; sampling from the distribution gives us a randomized 4/3-approximation algorithm. We note that known bad cases for the integrality gap have a gap of 4/3 and have a half-integral LP solution in which all the critical tight sets in the hierarchy are cycle cuts; thus our result is tight.

1 Introduction

In the traveling salesman problem (TSP), we are given a set of n cities and the costs c_{ij} of traveling from city i to city j for all i, j, and the goal of the problem is to find the least expensive tour that visits each city exactly once and returns to its starting point. An instance of the TSP is called *symmetric* if $c_{ij} = c_{ji}$ for all i, j. Costs obey the *triangle inequality* (or are *metric*) if $c_{ij} \leq c_{ik} + c_{kj}$ for all i, j, k. For ease of exposition, we consider the problem input as a complete graph $G = (V, E)$ for the set of cities V, with $c_e = c_{ij}$ for edge $e = (i, j)$. All instances we consider will be symmetric and obey the triangle inequality.

In a breakthrough result, Karlin, Klein, and Oveis Gharan [8] gave the first approximation algorithm with performance ratio better than 3/2, although the

A. Del Pia and V. Kaibel (Eds.): IPCO 2023, LNCS 13904, pp. 217–230, 2023.
https://doi.org/10.1007/978-3-031-32726-1_16

amount by which the bound was improved is quite small (approximately 10^{-36}). The algorithm follows the Christofides-Serdyukov template by selecting a random spanning tree from the max-entropy distribution, then using a T-join on the odd degree vertices of the tree to create a connected Eulerian subgraph.

One special case of the TSP is known as the *half-integral* case. To understand the half-integral case, we need to introduce a well-known LP relaxation of the TSP, sometimes called the *Subtour LP* or the *Held-Karp bound* [4,7], which is as follows:

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

where $\delta(S)$ is the set of all edges with exactly one endpoint in S and we use the shorthand that $x(F) = \sum_{e \in F} x_e$. A half-integral solution to the Subtour LP is one such that $x_e \in \{0, 1/2, 1\}$ for all $e \in E$, and a half-integer instance of the TSP is one whose LP solution is half-integral.

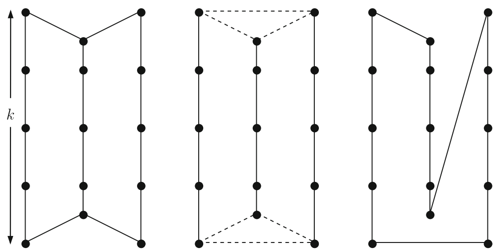

Fig. 1. Illustration of a known worst-case example for the integrality gap for the symmetric TSP with triangle inequality. The figure on the left gives an (unweighted) graph, and costs c_{ij} are the shortest path lengths in the graph. The figure in the center gives the LP solution, in which the dotted edges have value $1/2$, and the solid edges have value 1. The figure on the right gives the optimal tour. The ratio of the cost of the optimal tour to the value of the LP solution tends to $4/3$ as k increases.

The *integrality gap* of an LP relaxation is the worst-case ratio of an optimal integer solution to the linear program to the optimal linear programming solution. Wolsey [13] showed that the analysis of the Christofides-Seryukov algorithm could be used to show that the integrality gap of the Subtour LP is at most 3/2. It is known that the integrality gap of the Subtour LP is at least 4/3, due to a set of half-integral graph TSP instances shown in Fig. 1, and another set of half-integral weighted instances due to Boyd and Sebő [2] known as k-donuts. Schalekamp, Williamson, and van Zuylen [11] have conjectured that half-integral instances are the worst-case instances for the integrality gap. It has long been conjectured that the integrality gap is exactly 4/3, but until the work of Karlin et al. there had been no progress on the conjecture for several decades.

In the case of half-integral instances, some results are known. Mömke and Svensson [10] have shown a 4/3-approximation algorithm for half-integral graph TSP (in which cost c_{ij} is the number of edges in the shortest i-j path in an input graph), also yielding an integrality gap of 4/3 for such instances; because of the worst-case examples of Fig. 1, their result is tight. Boyd and Carr [1] give a 4/3-approximation algorithm (and an integrality gap of 4/3) for a subclass of half-integer solutions they call triangle points (in which the half-integer edges form disjoint triangles); the examples of Fig. 1 show that their result is tight also. Boyd and Sebő [2] give an upper bound of 10/7 for a subclass of half-integral solutions they call square points (in which the half-integer edges form disjoint 4-cycles). In a paper released just prior to their general improvement, Karlin, Klein, and Oveis Gharan [9] (KKO) gave a 1.49993-approximation algorithm in the half-integral case; in particular, they show that given a half-integral solution, they can produce a tour of cost at most 1.49993 times the value of the corresponding objective function. Gupta, Lee, Li, Mucha, Newman, and Sarkar [6] improve this factor to 1.4983.

With the improvements on the 3/2 bound remaining very incremental for weighted instances of the TSP, even in the half-integral case, we turn the question around and look for a large class of weighted half-integral instances for which we can prove that the 4/3 conjecture is correct, preferably one containing the known worst-case instances.

To define our instances, we turn to some terminology of KKO. The KKO result uses induction on a hierarchy of *critical tight* sets of the half-integral LP solution x. A set $S \subset V$ is *tight* if the corresponding LP constraint is met with equality; that is, $x(\delta(S)) = 2$. A set S is *critical* if it does not cross any other tight set; that is, for any other tight set T, either $S \cap T = \emptyset$ or $S \subseteq T$ or $T \subseteq S$. The critical tight sets then give rise to a natural tree-like hierarchy based on subset inclusion. KKO follow a Christofides-Serdyukov style algorithm that performs induction on the hierarchy. In their analysis, they differentiate between *cycle cuts* (in which the child nodes of a parent are linked by pairs of edges in a chain) and *degree cuts* (in which the child nodes of a parent form a 4-regular graph; more detail is given in subsequent sections).

In this paper, we will consider half-integral instances in which there are only cycle cuts, which we will refer to as half-integral cycle cut instances. Our contribution is to give a randomized $\frac{4}{3}$-approximation algorithm for these instances. More precisely, we give a distribution over connected Eulerian subgraphs such that each edge e is used with expectation at most $\frac{4}{3}x_e$, which implies the result (note that edges are sometimes doubled in the Eulerian graph). Our main theorem is as follows:

Theorem 1. *There is a randomized* $4/3$*-approximation algorithm for half-integral cycle cut instances of the TSP that produces an Eulerian tour with expected cost at most* $\frac{4}{3}\sum_{e\in E} c_e x_e$.

It is not hard to show that both the bad examples in Fig. 1 and the k-donut instances of Boyd and Sebő [2] are cycle cut instances (Boyd and Carr's result for triangle points works for the examples of Fig. 1, but not for k-donuts). Thus our bound of $4/3$ is tight and cannot be improved.

Our approach to the problem is novel and does not use the same Christofides-Serdyukov framework as employed by KKO and others. Instead, we perform a top-down induction on the hierarchy of critical tight sets. For each set in the hierarchy, we define a set of "patterns" of edges incident on it such that the set has even degree. For each pattern, we give a distribution of edges connecting the chain of child nodes in the cycle cut, which induces a distribution of patterns on each child. Crucially, we then show that there is a *feasible region* R of distributions over patterns, such that if the distribution of patterns on the parent node belongs to R, then the induced distribution on patterns on each child node also belongs to R. Our abstract is structured as follows. We give some needed preliminary definitions in Sect. 2. We then sketch our main result in Sect. 3, and conclude in Sect. 4. Due to space constraints, some proofs are omitted or sketched. The full paper can be accessed at https://arxiv.org/abs/2211.04639.

2 Preliminaries

Given a half-integral LP solution x, we construct a 4-regular 4-edge-connected multigraph $G = (V, E)$ by including a single copy of every edge e for which $x_e = \frac{1}{2}$ and two copies of every edge e for which $x_e = 1$. We state the following for general k-edge-connected multigraphs. In our setting, $k = 4$.

Definition 1. *For a* k*-edge-connected multigraph* $G = (V, E)$*, we say:*

- *Any set* $S \subseteq V$ *such that* $|\delta(S)| = k$ *(i.e., its boundary is a minimum cut) is a* ***tight set****.*
- *A set* $S \subseteq V$ *is* ***proper*** *if* $2 \leq |S| \leq n - 2$ *and a* ***singleton*** *if* $|S| = 1$.
- *Two sets* $S, S' \subseteq V$ ***cross*** *if all of* $S \setminus S'$*,* $S' \setminus S$*,* $S \cap S'$*, and* $V \setminus (S \cup S') \neq \emptyset$ *are non-empty.*

The following are two standard facts about minimum cuts; for proofs see [5].

Lemma 1. *If two tight sets* S *and* S' *cross, then each of* $S \setminus S'$*,* $S' \setminus S$*,* $S \cap S'$ *and* $\overline{S \cup S'}$ *are tight. Moreover, there are no edges from* $S \setminus S'$ *to* $S' \setminus S$*, and there are no edges from* $S \cap S'$ *to* $\overline{S \cup S'}$*.*

Lemma 2. *Let $G = (V, E)$ be a k-regular k-edge-connected graph. Suppose either $|V| = 3$ or G has at least one proper min cut, and every proper min cut is crossed by some other proper min cut. Then, k is even and G forms a cycle, with $k/2$ parallel edges between each adjacent pair of vertices.*

We now define our class of instances.

Definition 2 (Cycle cut instance). *We say a graph G is a* ***cycle cut instance*** *if every non-singleton tight set S can be written as the union of two tight sets $A, B \neq S$.*

As mentioned in the introduction this condition captures the two known integrality gap examples of the subtour LP.

We now show an equivalent definition of cycle cut instances after giving some definitions. First, fix an arbitrary **root vertex** $r \in V$, and for all cuts we consider we will take the side which does not contain r.

Definition 3 (Critical cuts). *A* critical *cut is any tight set $S \subseteq V \setminus \{r\}$ which does not cross any other tight set.*

Definition 4 (Hierarchy of critical cuts, $\mathcal{H}$). *Let $\mathcal{H} \subseteq 2^{V \setminus r}$ be the set of all critical cuts.*

The hierarchy naturally gives rise to a parent-child relationship between sets as follows:

Definition 5 (Child, parent, $E^{\rightarrow}(S)$). *Let $S \in \mathcal{H}$ such that $|S| \geq 2$. Call the maximal sets $C \in \mathcal{H}$ for which $C \subset S$ the children of S, and call S their parent. Finally, define $E^{\rightarrow}(S)$ to be the set of edges with endpoints in two different children of S.*

Definition 6 (Cycle cut, degree cut). *Let $S \in \mathcal{H}$ with $|S| \geq 2$. Then we call S a* ***cycle cut*** *if when $G \setminus S$ and all of the children of S are contracted, the resulting graph forms a cycle of length at least three with two parallel edges between each adjacent node. Otherwise, we call it a* ***degree cut****.*

While this definition of a cycle cut may sound specialized, due to Lemma 2, cycle cuts arise very naturally from collections of crossing min cuts.

Lemma 3. *If G is a cycle cut instance, then for any choice of r, $\mathcal{H}$ is composed only of cycle cuts (and singletons).*

One can also show that if for some choice of r, $\mathcal{H}$ is composed only of cycle cuts, then G is a cycle cut instance. Thus, in the remainder of the paper, we assume $\mathcal{H}$ is a collection of cycle cuts.

Given $S \in \mathcal{H}$, let $a_0 = G \setminus S$ and let $a_1, \ldots, a_k$ be its children in $\mathcal{H}$ (which are either vertices or cycle cuts). By Lemma 2 $a_0, \ldots, a_k$ can be arranged into a cycle such that two edges go between each adjacent vertex. WLOG let $a_1, \ldots, a_k$ be in counterclockwise order starting from a_0. We call a_1 the leftmost child of S and a_k the rightmost child.

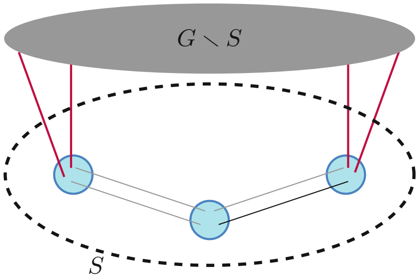

Fig. 2. S is an example of a cycle cut with three children. In blue are contracted critical tight sets. In gray is the rest of the graph with S contracted. As in Lemma 2, we can see that when $G \setminus S$ is contracted into a single vertex, the resulting graph is a cycle with 2 edges between each adjacent vertex. In our recursive proof of our main theorem in Sect. 3, we are given a distribution of Eulerian tours over G/S, so in particular on the red edges here, and will then extend it to G with the blue critical sets contracted by picking a distribution over the black edges. (Color figure online)

Definition 7 (External and internal cycles cuts). *Let $S \in \mathcal{H}$ such that $S \neq V \setminus \{r\}$ be a cut with parent S'. We call S external if in the ordering $a_0, \dots, a_k$ of S' (as given above), $S = a_1$ or $S = a_k$. Otherwise, call S internal.*

For example, if the blue nodes in Fig. 2 are contracted cycle cuts, the left and right nodes are external, while the middle one is internal. Note that for an cycle cut S with parent S', if S is external then $|\delta(S) \cap \delta(S')| = 2$, and if S is internal then $|\delta(S) \cap \delta(S')| = 0$.

Using the following simple fact, we will now describe our convention for drawing and describing cycle cuts:

Lemma 4. *Let $A, B, C \in \mathcal{H}$ be three distinct critical cuts such that $A \subsetneq B$ and $B \cap C = \emptyset$ or $B \subseteq C$. Then $|\delta(A) \cap \delta(C)| \leq 1$.*

Definition 8 ($\delta^L(S), \delta^R(S)$). *Let $S \in \mathcal{H}$ be a cycle cut. We will define a partition of $\delta(S)$ into two sets $\delta^L(S), \delta^R(S)$ each consisting of two edges.*

If $S \neq V \setminus \{r\}$, then it has a parent S'. S' has children $a_1, \dots, a_k$ such that $S = a_i$ for $i \neq 0$. Let $\delta^L(S) = \delta(S) \cap \delta(a_{i-1})$ and $\delta^R(S) = \delta(S) \cap \delta(a_{i+1 \pmod{k+1}})$. In other words, we partition the edges of S into the two edges going to the left neighbor of S in the cycle defined by S''s children and the two edges going to the right neighbor.

Otherwise $S = V \setminus \{r\}$. Then if $a_1, \dots, a_k$ are the children of S, let $\delta^L(S)$ consist of an arbitrary edge from $\delta(a_1) \cap \delta(S)$ and an arbitrary edge from $\delta(a_k) \cap \delta(S)$. Let $\delta^R(S) = \delta(S) \setminus \delta^L(S)$.

By Lemma 4 and the definition of $\delta^L(S), \delta^R(S)$ for $S = V \setminus \{r\}$, if S' is an external child of a cycle cut S, then $|\delta^L(S) \cap \delta(S')| = |\delta^R(S) \cap \delta(S')| = 1$. This allows us to adopt the following convention for drawing cycle cuts which we will call the **caterpillar drawing** of S: for an example, see Fig. 3. Formally, let $S \in \mathcal{H}$ be a cycle cut with children $a_1, \ldots, a_k \in \mathcal{H}$. Arrange $a_1, \ldots, a_k$ in a horizontal line. First, expand a_1 vertically into its children (if it is not a singleton) such that the unique edge in $\delta^L(S) \cap \delta(a_1)$ is pointing up (if it is a singleton, simply draw this edge pointing up. Then, expand $a_2, \ldots, a_k$ one by one into their respective children (if they exist), placing the children vertically in increasing or decreasing order of their index so that the edges from a_i to a_{i+1} do not cross. If a_k is a singleton, arbitrarily choose which edge to draw pointing up. Otherwise, let a' be the topmost child of a_k. Draw the unique edge in $\delta(S) \cap \delta(a')$ pointing up. There are two types of cycle cuts:

Definition 9 (Straight and twisted cycle cuts). *Let $S \in \mathcal{H}$ be a cycle cut. If $\delta^L(S)$ has both edges pointing up in the caterpillar drawing of S, then call it a straight cycle cut. Otherwise, call it a twisted cycle cut. See Fig. 3 for examples.*

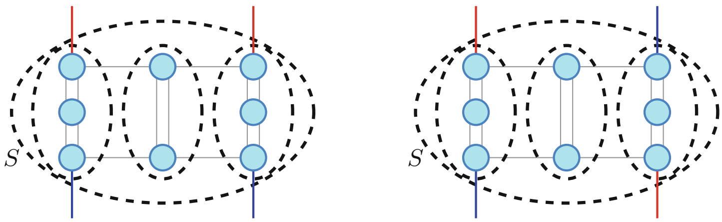

Fig. 3. Caterpillar drawings of two different cycle cuts S. The red edges are in the $\delta^L(S)$ partition, and the blue edges are in the $\delta^R(S)$ partition. The left drawing is a straight cycle cut, and the right is a twisted cycle cut as per Definition 9 (Color figure online).

In the next section, we abbreviate the caterpillar drawing by contracting the non-singleton children of S (see Fig. 4). We do so partially for cleaner pictures but also to emphasize that all the relevant information used by our construction in the following section is contained in the abbreviated pictures.

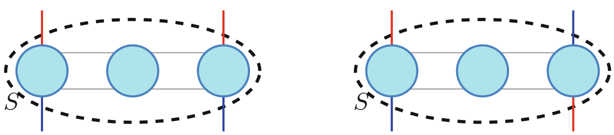

Fig. 4. On the left is a shorthand caterpillar drawing for the straight cycle cut on the left in Fig. 3 obtained by contracting its children. Similarly for the right. We will use this style of picture in future sections.

3 Proof of Theorem 1

We now present a summary of the proof of our main result, a $\frac{4}{3}$-approximation for half-integral cycle-cut instances of the TSP. To prove Theorem 1, we construct a distribution of Eulerian tours such that every edge is used at most $\frac{2}{3}$ of the time. Since $x_e = \frac{1}{2}$ for every edge in the graph, this immediately implies that when we sample a tour from this distribution, its expected cost is at most $\frac{4}{3}$ times the value of the LP. We work on the cycle cut hierarchy from the top down, and inductively specify the distribution of edges that enter every cut.

Figure 4 depicts our convention for visualizing a cycle cut as described in Sect. 2. We say that a cycle cut is *even* if it contains an even number of children, and *odd* otherwise. Fig. 6 illustrates the *patterns* we use, where "pattern" refers to a multiset of edges that enter a cycle cut. For each pattern entering a parent cycle cut, we give (randomized) rules which describe how to connect up its children – this induces a distribution of patterns entering each child. We represent this process using a Markov chain with 4 states, illustrated in Fig. 6. The figure shows the mapping from patterns to states; the transitions will come from the rules for connecting up the children, which we describe later. In the figure, each state contains two pictures, which represent the *parity* of the edges in the patterns that are mapped to the state. Specifically, a present edge is used exactly once, whereas an edge that is not present may be either unused or doubled. For example, Fig. 7 illustrates all possible patterns that are captured by the top picture of state 1. Finally, we maintain the invariant that if a cycle cut is in a given state, then each of the two pictures are equally likely. (When we later give the rules for connecting up the children, we will ensure this invariant is preserved.) Thus, when we say a cycle cut is in a given state with probability p, this means the parity of the pattern entering it follows the top picture in the state with probability $\frac{p}{2}$, and the bottom picture with probability $\frac{p}{2}$. We will use the phrase "the distribution of patterns entering a cycle cut C is (p_1, p_2, p_3, p_4)" to mean that for all $i \in \{1, 2, 3, 4\}$, C is in state i with probability p_i.

To prove our main result, we will give a *feasible region* R of distributions over the states of the Markov chain, such that: 1) If the distribution of patterns entering a cycle cut C belongs to R, there is a way to connect up the children of C such that the distribution on each child also belongs to R, and 2) for each $\mathbf{p} \in R$, the corresponding rule for connecting the children of C uses each edge in $E^{\rightarrow}(C)$ at most $\frac{2}{3} = \frac{4}{3}x_e$ of the time in expectation. The feasible region is given in Definition 10. As long as R is nonempty, 1) and 2) are sufficient to give the result since we can induce any distribution on the cycle cut $V \setminus \{r\}$.

Definition 10 (The Feasible Region). *Let*

$$R = \left\{(p_1, p_2, p_3, p_4) \in \mathbb{R}^4_+ : p_1 + p_2 + p_3 + p_4 = 1,\ p_1 + p_2 = \frac{2}{3},\ p_2 + p_4 \geq \frac{1}{3}\right\}.$$

See Fig. 5 for an visualization of R in a 2-dimensional space.

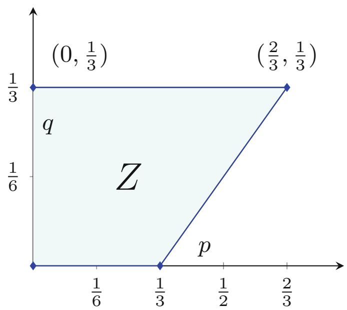

Fig. 5. The feasible region of distributions is $R = \{(p, \frac{2}{3} - p, \frac{1}{3} - q, q) : (p, q) \in Z\}$, where Z is the polytope above.

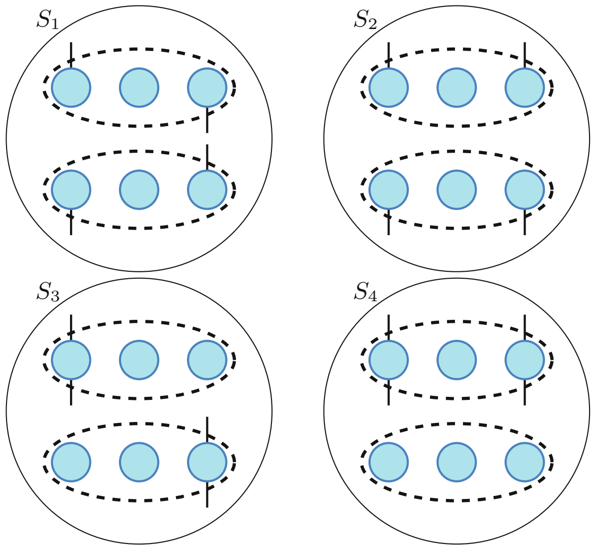

Fig. 6. The patterns and how they map to states of a Markov chain. The states are unchanged regardless of the number of children: they are defined only with respect to which of the edges are in. Note that we ignore doubled edges.

To describe the transitions of the Markov chain, we give (randomized) rules that dictate, for a cycle cut C and a pattern entering it, how to connect up its children. These rules depend on whether C is even or odd. The final form of the Markov chains is illustrated in Fig. 8.[1] The meaning of taking one transition is as follows. Suppose the distribution of patterns entering C is (p_1, p_2, p_3, p_4),

[1] In the figure, if there is a variable on an arc, it means that any transition probability in the range of that variable is possible. For example, in P_{even}, we can transition from S_2 to S_1 with probability z for any $z \in [0, 1]$; the transition from S_2 to S_3 then happens with probability $1 - z$.

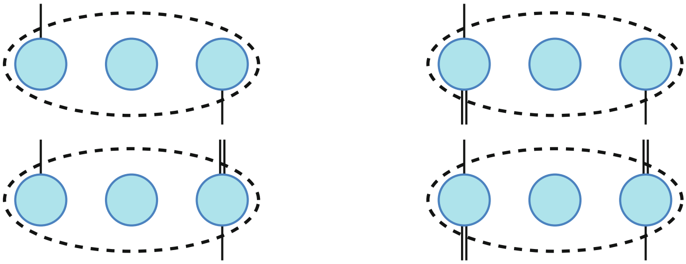

Fig. 7. In our illustrations of the patterns entering a given cycle cut, any edge that is not present may either be unused or doubled. Therefore, all four of the given edge configurations are represented by the upper left most state, S_1.

and suppose (q_1, q_2, q_3, q_4) is the resulting distribution after one transition of a Markov chain. What this means is that for each child of C, the distribution of patterns entering it will be **either** (q_1, q_2, q_3, q_4) or (q_2, q_1, q_3, q_4) depending on if the child is straight or twisted, respectively (see Definition 9 and Fig. 3). In particular, it can be shown that if (q_1, q_2, q_3, q_4) is the distribution induced on a child which is a straight cycle cut, then (q_2, q_1, q_3, q_4) would be the distribution induced on a child which is a twisted cycle cut. Thus, it is sufficient to check that: i) the distributions induced on straight children lie in the feasible region and ii) if (q_1, q_2, q_3, q_4) is a distribution induced on straight children, then (q_2, q_1, q_3, q_4) is also in the feasible region. This corresponds to the set of distributions induced on the children being symmetric under this transformation.[2]

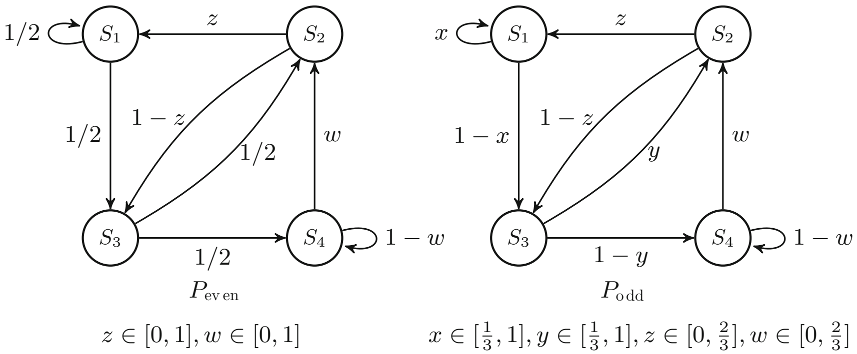

Fig. 8. The variables on the arcs indicate that one can feasibly transition according to any probability in the range.

[2] Note that the feasible region is not symmetric under this transformation. The distribution induced on the children is thus a symmetric subset of the feasible region.

Proposition 1. *For any cycle cut $C \in \mathcal{H}$ and any distribution of patterns entering C, there is a way to connect its children so that the induced distribution on each child is given by 1) applying the corresponding Markov chain in Fig. 8, and then 2) swapping the first two coordinates if the child is twisted.*

Proof (Sketch). The proof involves going through the 8 cases one by one (depending on the parity of the cut, and which of the 4 states it is in), and showing that in each case, there is a (randomized) rule for connecting the children that achieve the transitions in Fig. 8. To illustrate the main idea, we show the rule in the case that C is even and in state 4.

In this case, the rule for connecting the children of C is illustrated in Fig. 9. Let $w \in [0, 1]$. With probability w, we make all children transition to state 2. To do this, first suppose C has all 4 single edges entering it (the top picture in the left box). In this case, we consider the pairs of edges in $E^{\rightarrow}(C)$ from left to right, and alternate 1) doubling one of the two edges with equal probability (shown by the dotted black edges), and 2) using both edges (shown by the solid black edges). Because C is even, the rightmost pair of edges ends up falling in case 1) of the alternating rule, and so all children transition to state 2. The case where all the edges entering C are used an even number of times (the bottom picture in the left box) is quite similar, except we begin the alternating rule by using both edges.

On the other hand, with probability $1 - w$, we transition back to state 4. This is accomplished by using each pair of edges in the top case of state 4, and by doubling one edge from each pair uniformly at random in the bottom case of state 4. The net transition probabilities are then $(0, w, 0, 1-w)$, where w can be any number from 0 to 1. $\square$

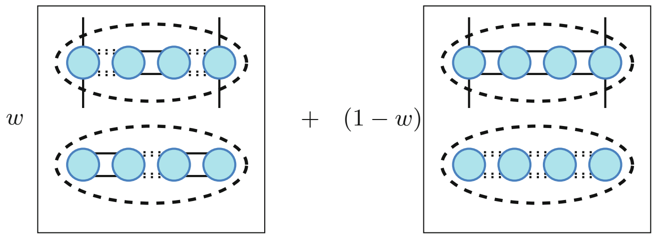

Fig. 9. Transition for state 4 in the even case.

We ensure that in all cases, each edge in $E^{\rightarrow}(C)$ is used $\frac{1}{2}$, $\frac{1}{2}$, 1, 1 times in expectation if the pattern entering C belongs to state 1, 2, 3, 4, respectively. Therefore, if $\mathbf{p} = (p_1, p_2, p_3, p_4)$ are the probabilities that we are in states 1, 2, 3, 4 respectively, then each edge in $E^{\rightarrow}(C)$ is used exactly $\frac{1}{2}p_1 + \frac{1}{2}p_2 + p_3 + p_4 = 1 - \frac{1}{2}(p_1 + p_2)$ of the time in expectation. Thus to get a $\frac{4}{3}$-approximation, it is necessary that $p_1 + p_2 \geq \frac{2}{3}$. Note that if $\mathbf{p} \in R$, then $p_1 + p_2 = \frac{2}{3}$, so that each edge is used exactly $\frac{2}{3}$ of the time.

To complete the proof, we only need show that if the distribution of patterns entering a cycle cut C belongs to R, then the induced distributions on the children also belong to R. Thus R is sufficient, in sense that if the distribution entering a cycle cut belongs to R, then it is possible to get a $\frac{4}{3}$-approximation all the way down the hierarchy using the Markov chains in Fig. 8. Moreover, we are able to show that R is *necessary*; if the distribution entering a cycle cut does **not** belong to R, then it is impossible to obtain a $\frac{4}{3}$-approximation using our Markov chains. In this sense, R is the *largest* feasible region using our technique.

Theorem 2. *1. (R is **sufficient**) If the distribution of patterns entering a cycle cut belongs to R, then there are feasible Markov chains (among the ones shown in Fig. 8) such that the induced distribution entering each child also belongs to R.*

*2. (R is **necessary**) Suppose the distribution of patterns entering a cycle cut does **not** belong to R. Then it is not possible to obtain a $\frac{4}{3}$-approximation using the Markov chains in Fig. 8.*

Proof (Sketch). For 1), we show that for any $\mathbf{p} \in R$ and for C even or odd, there are feasible values for the transition probabilities of the corresponding Markov chain such that the resulting distribution $\mathbf{q} \in R$ (and also $\mathbf{q}$ with its first two coordinates swapped is in R.) The values of the transition probabilities are derived as a function of $\mathbf{p}$. For 2), we consider an arbitrary distribution $\mathbf{p}$ (not necessarily in R), and let $\mathbf{q}^{(1)}$ and $\mathbf{q}^{(2)}$ be the distributions obtained by applying P_{even} once and twice, respectively. We then argue that $\mathbf{p}$ must belong to R in order for $\mathbf{q}^{(1)}$ and $\mathbf{q}^{(2)}$ to each have their first two coordinates sum to at least $\frac{2}{3}$. □

Example. To give the reader some more intuition, we give a specific example of how to maintain distributions in R on all the cuts in the hierarchy by choosing appropriate transition probabilities on the Markov chains in Fig. 8. Let $\mathbf{p} = (\frac{4}{9}, \frac{2}{9}, \frac{2}{9}, \frac{1}{9})$ and $\mathbf{q} = (\frac{2}{9}, \frac{4}{9}, \frac{2}{9}, \frac{1}{9})$ (i.e. $\mathbf{q}$ is $\mathbf{p}$ with the first two coordinates swapped). It is easy to check that $\mathbf{p}, \mathbf{q} \in R$. We now show for any half-integral cycle cut instance, it is possible to make it so that the distribution entering any cycle cut is either $\mathbf{p}$ or $\mathbf{q}$.

To see this, let C be a cycle cut and suppose C is odd. Set the transition probabilities in P_{odd} to be $x = y = z = w = \frac{2}{3}$. For these probabilities, it is easy to check that $P_{\mathsf{odd}}\mathbf{p} = P_{\mathsf{odd}}\mathbf{q} = \mathbf{p}$.[3] On the other hand, if C is even, setting $z = w = 1$ in P_{even} gives $P_{\mathsf{even}}\mathbf{p} = \mathbf{p}$, and setting $z = \frac{3}{4}, w = 1$ gives $P_{\mathsf{even}}\mathbf{q} = \mathbf{p}$. Thus, as long as the distribution entering C is $\mathbf{p}$ or $\mathbf{q}$, we can make the distribution on each child of C be either $\mathbf{p}$ or $\mathbf{q}$.

Together with Proposition 1, this already proves a $\frac{4}{3}$-approximation for half-integral cycle cut instances. The additional contribution of Theorem 2 is an exact characterization of the region of distributions that give a $\frac{4}{3}$-approximation using our techniques.

---

[3] In fact, it can be checked that for these probabilities, P_{odd} maps *every* distribution (whose first two coordinates sum to $\frac{2}{3}$), to $\mathbf{p}$.

4 Conclusion and Open Questions

Our result leads to several interesting open questions. One such open question is whether our result extends to the case of cycle cuts for non-half-integral solutions. We believe this to be possible through a more refined understanding of the patterns that result from considering non-half-integral solutions.

Clearly a better understanding of what happens in the case of degree cuts is needed to make substantial progress on the overall half-integral case. We think it is possible to improve incrementally on the 1.4983-approximation of Gupta et al. [6] by using a combination of ideas from this paper with a few other small improvements. Recall that in a degree cut, each vertex has degree four, there are no parallel edges, and every proper cut has at least six edges crossing it. Ideally one would be able to show that any distribution on a parent cut lying in the feasible region of Fig. 5 could be used to induce a distribution on patterns of the children of the degree cut in a subregion of the feasible region with each edge used at most $2/3$ of the time; such a result would lead immediately to a $4/3$ integrality gap for half-integral instances.

Acknowledgment. The first and third authors would like to thank Anke van Zuylen for early discussions on this problem. The first and third authors were supported in part by NSF grant CCF-2007009. The first author was also supported by NSERC fellowship PGSD3-532673-2019. The second author was supported in part by NSF grants DGE-1762114, CCF-1813135, and CCF-1552097. We would like to thank Martin Drees for his helpful suggestions that allowed us to simplify the proof of the main result.

References

1. Boyd, S., Carr, R.: Finding low cost TSP and 2-matching solutions using certain half-integer subtour vertices. Discret. Optim. **8**, 525–539 (2011)
2. Boyd, S., Sebő, A.: The salesman's improved tours for fundamental classes. Math. Program. **186**, 289–307 (2021)
3. Christofides, N.: 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)
4. Dantzig, G., Fulkerson, R., Johnson, S.: Solution of a large-scale traveling-salesman problem. J. Oper. Res. Soc. Am. **2**(4), 393–410 (1954). ISSN 00963984. URL https://www.jstor.org/stable/166695
5. Fleiner, T., Frank, A.: A quick proof for the cactus representation of mincuts. Technical Report QP-2009-03, Egerváry Research Group, Budapest (2009). https://www.cs.elte.hu/egres
6. Gupta, A., Lee, E., Li, J., Mucha, M., Newman, H., Sarkar, S.: Matroid-based TSP rounding for half-integral solutions. In: Aardal, K., Sanità, L. (eds.) Integer Programming and Combinatorial Optimization. LNCS, vol. 13265, pp. 305–318 (2022). https://doi.org/10.1007/978-3-031-06901-7_23,See also https://arxiv.org/pdf/2111.09290.pdf
7. Held, M., Karp, R.M.: The traveling-salesman problem and minimum spanning trees. Oper. Res. **18**, 1138–1162 (1971)

8. Karlin, A.R., Klein, N., Gharan, S.O.: A (slightly) improved approximation algorithm for metric tsp. In: STOC. ACM (2021)
9. Karlin, A.R., Klein, N., Gharan, S.O.: An improved approximation algorithm for TSP in the half integral case. In: Makarychev, K., Makarychev, Y., Tulsiani, M., Kamath, G., Chuzhoy, J. (eds.) STOC, pp. 28–39. ACM (2020)
10. Mömke, T., Svensson, O.: Removing and adding edges for the traveling salesman problem. J. ACM, 63 (2016). Article 2
11. 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)
12. Serdyukov, A.: On some extremal walks in graphs. Upravlyaemye Sistemy **17**, 76–79 (1978)
13. Wolsey, L.A.: Heuristic analysis, linear programming and branch and bound. Math. Program. Study **13**, 121–134 (1980)