

An Improved Integrality Gap for Asymmetric TSP Paths*

Zachary Friggstad[†]

Anupam Gupta[‡]

Mohit Singh[§]

February 14, 2013

Abstract

The Asymmetric Traveling Salesperson Path (ATSP) problem is one where, given an *asymmetric* metric space (V, d) with specified vertices s and t , the goal is to find an s - t path of minimum length that passes through all the vertices in V .

This problem is closely related to the Asymmetric TSP (ATSP) problem, which seeks to find a tour (instead of an s - t path) visiting all the nodes: for ATSP, a ρ -approximation guarantee implies an $O(\rho)$ -approximation for ATSP. However, no such connection is known for the *integrality gaps* of the linear programming relaxations for these problems: the current-best approximation algorithm for ATSP is $O(\log n / \log \log n)$, whereas the best bound on the integrality gap of the natural LP relaxation (the subtour elimination LP) for ATSP is $O(\log n)$.

In this paper, we close this gap, and improve the current best bound on the integrality gap from $O(\log n)$ to $O(\log n / \log \log n)$. The resulting algorithm uses the structure of narrow s - t cuts in the LP solution to construct a (random) tree witnessing this integrality gap. We also give a simpler family of instances showing the integrality gap of this LP is at least 2.

1 Introduction

In the Asymmetric Traveling Salesperson Path (ATSP) problem, we are given an *asymmetric* metric space (V, d) (i.e., one where the distances satisfy the triangle inequality, but potentially not the symmetry condition), and also specified source and sink vertices s and t , and the goal is to find an s - t Hamilton path of minimum length.

This ATSP problem is a close relative of the Asymmetric TSP problem (ATSP), where the goal is to find a Hamilton tour instead of an s - t path. For this ATSP problem, the $\log_2 n$ -approximation of Frieze, Galbiati, and Maffioli [9] from 1982 was the best result known for more than two decades, until it was finally improved by constant factors in [4, 11, 8]. A breakthrough on this problem was an $O(\frac{\log n}{\log \log n})$ -approximation result due to Asadpour, Goemans, Mądry, Oveis Gharan, and Saberi [2]; they also bounded the integrality gap of the subtour elimination linear programming relaxation for ATSP by the same factor.

*An extended abstract of this paper appears in the *Proceedings of the 16th Conference on Integer Programming and Combinatorial Optimization, 2013*.

[†]Department of Combinatorics and Optimization, University of Waterloo.

[‡]Department of Computer Science, Carnegie Mellon University, Pittsburgh PA 15213, and Microsoft Research SVC, Mountain View, CA 94043. Research was partly supported by NSF awards CCF-0964474 and CCF-1016799.

[§]Microsoft Research, Redmond.

Somewhat surprisingly, the study of ATSP has been of a more recent vintage: the first approximation algorithms appeared only around 2005 [13, 6, 8]. It is easily seen that the ATSP reduces to ATSP in an approximation preserving fashion (by guessing two consecutive nodes on the tour). In the other direction, [8] showed that a ρ -approximation to the ATSP problem implies an $O(\rho)$ -approximation to the ATSP problem. Using the above-mentioned $O(\frac{\log n}{\log \log n})$ -approximation for ATSP [2], this implies an $O(\frac{\log n}{\log \log n})$ -approximation for ATSP as well.

The subtour elimination linear program generalizes simply to the ATSP problem and is given in Section 2. However, the best previous integrality gap for this LP for ATSP was $O(\log n)$ [10]. In this paper we show the following result.

Theorem 1.1. *The integrality gap of the subtour elimination linear program for the ATSP problem is at most $O(\frac{\log n}{\log \log n})$.*

We also give a simple construction showing that the integrality gap of this LP is at least 2; this example is simpler than previous known integrality gap instance showing the same lower bound, due to Charikar, Goemans, and Karloff [5].

Given the central nature of linear programs in approximation algorithms, it is useful to understand the integrality gaps for linear programming relaxations of optimization problems. Not only does this study give us a deeper understanding into the underlying problems, but also upper bounds on the integrality gap of LPs are often required for some reductions to go through. For example, the polylogarithmic approximation guarantees in the work of Nagarajan and Ravi [14] for Directed Orienteering and Minimum Ratio Rooted Cycle, and those in the work of Bateni and Chuzhoy [3] for Directed k -Stroll and Directed k -Tour were all improved by a factor of $\log \log n$ following the improved bound of $O(\frac{\log n}{\log \log n})$ on the integrality gap of the subtour LP relaxation for ATSP. Note that these improvements do not follow merely from improved approximation guarantees.

1.1 Our Approach

Our approach to bound the integrality gap for ATSP is similar to that for ATSP [2], but with some crucial differences. We sample a random spanning tree and then augment the directed version of this tree to an integral circulation using Hoffman’s circulation theorem while ensuring the t - s edge is only used once. Following the corresponding Eulerian circuit and deleting the t - s edge results in a spanning s - t walk.

However, the non-Eulerian nature of the ATSP problem makes it difficult to satisfy the cut requirements in Hoffman’s circulation theorem if we sample the spanning tree directly from the distribution given by the LP solution. It turns out that the problems come from the s - t cuts U that are nearly-tight: i.e., which satisfy $1 < x^*(\partial^+(U)) < 1 + \tau$ for some small constant τ — these give rise to problems when the sampled spanning tree includes more than one edge across this cut. Such problems also arise in the symmetric TSP paths case (studied in a recent paper of An, Kleinberg, and Shmoys [1]): their approach is again to take a random tree directly from the distribution given by the optimal LP solution, but in some cases they need to boost the narrow cuts, and they show that the loss due to this boosting is small.

In our case, the asymmetry in the problem means that boosting the narrow cuts might be prohibitively expensive. Hence, our idea is to preprocess the distribution given by the LP solution to *tighten* the narrow cuts, so that we never pick two edges from a narrow cut. Since the original LP solution lies in the spanning tree polytope, lowering the solution on some edges means we need to raise other edges, which causes the costs to increase, and technical heart of the paper is to ensure this can be done with little extra loss.

1.2 Other Related Work

The first non-trivial approximation for ATSP was an $O(\sqrt{n})$ -approximation by Lam and Newman [13]. This was improved to $O(\log n)$ by Chekuri and Pál [6], and the constant was further improved in [8]. The paper [8] also showed that ATSP and ATSP had approximability within a constant factor of each other. All these results are combinatorial and do not bound integrality gap of ATSP. A bound of $O(\sqrt{n})$ on the integrality gap of ATSP was given by Nagarajan and Ravi [15], and was improved to $O(\log n)$ by Friggstad, Salavatipour and Svitkina [10]. Note that there is still no result known that relates the integrality gaps of the ATSP and ATSP problems in a black-box fashion.

In the symmetric case (where the problems become TSPP and TSP respectively), constant factor approximations and integrality gaps have long been known. We do not survey the rich body of literature on TSP here, instead pointing the reader to, e.g., the recent paper on graphical TSP by Sebö and Vygen [17]. It is, however, important to mention the recent 1.618-approximation for TSPP in a beautiful new result by An, Kleinberg, and Shmoys [1]. They proceed via bounding the integrality gap of the LP relaxation, and their algorithm also proceeds via studying the narrow s - t cuts; the connections to their work are discussed in Section 1.1.

1.3 Notation and Preliminaries

Given a directed graph $G = (V, A)$, and two disjoint sets $U, U' \subseteq V$, let $\partial(U; U') = A \cap (U \times U')$. We use the standard shorthand that $\partial^+(U) := \partial(U; V \setminus U)$, and $\partial^-(U) := \partial(V \setminus U; U)$. When the set U is a singleton (say $U = \{u\}$), we use $\partial^+(u)$ or $\partial^-(u)$ instead of $\partial^+(\{u\})$ or $\partial^-(\{u\})$. For undirected graph $H = (V, E)$, we use $\partial(U; U')$ to denote edges crossing between U and U' , and $\partial(U)$ to denote the edges with exactly one endpoint in U (which is the same as $\partial(V \setminus U)$).

For a digraph $G = (V, A)$, a set of arcs $B \subseteq A$ is *weakly connected* if the undirected version of B forms a connected graph that spans all vertices in A .

For values $x_a \in \mathbb{R}$ for all $a \in A$, and a set of arcs $B \subseteq A$, we let $x(B)$ denote the sum $\sum_{a \in B} x_a$.

Given an undirected graph $H = (V, E)$, we let $\chi_T \in \{0, 1\}^{|E|}$ denote the characteristic vector of a spanning tree T , then the spanning tree polytope is the convex hull of $\{\chi_T \mid T \text{ spanning tree of } H\}$. See, e.g., [16, Chapter 50] for several equivalent linear programming formulations of this polytope. We sometimes abuse notation and call a set of directed arcs T a tree if the undirected version of T is a tree in the usual sense.

An *directed metric graph* on vertices V has arcs $A = \binom{V}{2}$ where the non-negative arc costs satisfy the triangle inequality $c_{uv} + c_{vw} \geq c_{uw}$ for all $u, v, w \in V$. However, arcs uv and vu need not have the same cost. An instance of the ATSP problem is a directed metric graph along with distinguished vertices $s \neq t$.

2 The Rounding Algorithm

In this section, we give the linear programming relaxation for the Asymmetric TSP Path problem, and show how to round it to get a path of cost at most $O(\frac{\log n}{\log \log n})$ times the cost of the optimal LP solution. We then give the proof, with some of the details being deferred to the following sections.

Given a directed metric graph $G = (V, A)$ with arc costs $\{c_a\}_{a \in A}$, we use the following standard linear

programming relaxation for ATSP which is also known as the subtour elimination linear program.

$$\begin{aligned}
& \text{minimize : } \sum_{a \in E} c_a x_a && \text{(ATSP)} \\
& \text{s.t. : } x(\partial^+(s)) = x(\partial^-(t)) = 1 && (1) \\
& x(\partial^-(s)) = x(\partial^+(t)) = 0 && (2) \\
& x(\partial^+(v)) = x(\partial^-(v)) = 1 && \forall v \in V \setminus \{s, t\} \quad (3) \\
& x(\partial^+(U)) \geq 1 && \forall \{s\} \subseteq U \subsetneq V \quad (4) \\
& x_a \geq 0 && \forall a \in E
\end{aligned}$$

We begin by solving the above LP to obtain an optimal solution x^* . Consider the undirected (multi)graph $H = (V, E)$ obtained by removing the orientation of the arcs of G . That is, create precisely two edges between every two nodes $u, v \in V$ in H , one having cost c_{uv} and the other having cost c_{vu} . (Hence, $|E| = |A|$.) For a point $w \in \mathbb{R}_+^A$, let $\kappa(w)$ denote the corresponding point in \mathbb{R}_+^E , and view $\kappa(w)$ as the “undirected” version of w .

We will use the following definition: An s - t cut is a subset $U \subset V$ such that $\{s\} \subseteq U \subseteq V \setminus \{t\}$. The LP constraints imply that $x^*(\partial^+(U)) - x^*(\partial^-(U)) = 1$ for every s - t cut U . Also, $x^*(\partial^+(U)) = x^*(\partial^-(U)) \geq 1$ for every nonempty $U \subseteq V \setminus \{s, t\}$.

Definition 2.1 (Narrow cuts). Let $\tau \geq 0$. An s - t cut U is τ -narrow if $x^*(\partial^+(U)) < 1 + \tau$ (or equivalently, $x^*(\partial^-(U)) < \tau$).

The main technical lemma is the following:

Lemma 2.2. For any $\tau \in [0, 1/4]$, one can find, in polynomial-time, a vector $z \in [0, 1]^A$ (over the directed arcs) such that:

- (a) its undirected version $\kappa(z)$ lies in the spanning tree polytope for H ,
- (b) $z \leq \frac{1}{1-3\tau} x^*$ (where the inequality denotes component-wise dominance), and
- (c) $z(\partial^+(U)) = 1$ and $z(\partial^-(U)) = 0$ for every τ -narrow s - t cut U .

Before we prove the lemma (in Section 2.1), let us sketch how it will be useful to get the ATSP. Since z (or more correctly, its undirected version $\kappa(z)$) lies in the spanning tree polytope, it can be represented as a convex combination of spanning trees. Using some recently-developed algorithms (e.g., those due to [2, 7]) one can choose a spanning tree that crosses each cut only $O(\frac{\log n}{\log \log n})$ times more than the LP solution. Finally, we can use $O(\frac{\log n}{\log \log n})$ times the LP solution to patch this tree to get an s - t path. Since the LP solution is “weak” on the narrow cuts and may contribute very little to this patching (at most τ), it is crucial that by property (c) above, this tree will cross the narrow cuts *only once*, and that too, it crosses in the “right” direction, so we never need to use the LP when verifying the cut conditions of Hoffman’s circulation theorem on narrow cuts. The details of these operations appear in Section 3.

2.1 The Structure of Narrow Cuts

We now prove Lemma 2.2: it says that we can take the LP solution x^* and find another vector z such that if a s - t cut is narrow in x^* (i.e., the total x^* value crossing the cut lies in $[1, 1 + \tau]$), then z crosses it to an extent

precisely 1. Moreover, the undirected version of z can be written as a convex combination of spanning trees, and z_a is not much larger than x_a^* for any arc a .

Note that the undirected version of x^* itself can be written as a convex combination of spanning trees, so if we force z to cross the narrow cuts to an extent less than x^* (loosely, this reduces the connectivity), we'd better increase the value on other arcs. To show we can perform this operation without changing any of the coordinates by very much, we need to study the structure of narrow cuts more closely. (Such a study is done in the *symmetric* TSP path paper of An et al. [1], but our goals and theorems are somewhat different.)

First, say two s - t cuts U and W *cross* if $U \setminus W$ and $W \setminus U$ are non-empty.

Lemma 2.3. *For $\tau \leq 1/4$, no two τ -narrow s - t cuts cross.*

Proof. Suppose U and W are crossing τ -narrow s - t cuts. Then

$$\begin{aligned} 2 + 2\tau &> x^*(\partial^+(U)) + x^*(\partial^+(W)) \\ &= x^*(\partial^+(U \setminus W)) + x^*(\partial^+(W \setminus U)) + x^*(\partial^+(U \cap W)) \\ &\quad + x^*(\partial(U \cap W; V \setminus (U \cup W))) - x^*(\partial(U \cup W) \setminus (U \cap W); U \cap W) \\ &\geq 1 + 1 + 1 + 0 - 2\tau \\ &= 3 - 2\tau \end{aligned}$$

where the last inequality follows from the first three terms being cuts excluding t and hence having at least unit x^* -value crossing them (by the LP constraints), the fourth term being non-negative, and the last term being the x^* -value of subset of the arcs in $\partial^-(U) \cup \partial^-(W)$ and remembering that U and W are τ -narrow. However, this contradicts $\tau \leq 1/4$. \square

Lemma 2.3 says that the τ -narrow cuts form a chain $\{s\} = U_1 \subset U_2 \subset \dots \subset U_k = V \setminus \{t\}$ with $k \geq 2$. For $1 < i \leq k$, let $L_i := U_i \setminus U_{i-1}$. We also define $L_1 = \{s\}$ and $L_{k+1} = \{t\}$. Let $L_{\leq i} := \bigcup_{j=1}^i L_j$ and $L_{\geq i} := \bigcup_{j=i}^{k+1} L_j$. For the rest of this paper, we will use τ to denote a value in the range $[0, 1/4]$. Ultimately, we will set $\tau := 1/4$ for the final bound but we state the lemmas in their full generality for $\tau \leq 1/4$.

Next, we show that out of the (at most) $1 + \tau$ mass of x^* across each τ -narrow cut U_i , most of it comes from the ‘‘local’’ arcs in $\partial(L_i; L_{i+1})$.

Lemma 2.4. *For each $1 \leq i \leq k$; $x^*(\partial(L_i; L_{i+1})) \geq 1 - 3\tau$.*

Proof. For $i = 1$, since $s, t \notin L_2$ then we have $x^*(\partial^-(L_2)) \geq 1$ from the LP constraints. We also have $x^*(\partial^-(U_2)) < \tau$ since U_2 is τ -narrow, so at least $1 - \tau$ of $x^*(\partial^-(L_2))$ comes from $x^*(\partial(L_1; L_2))$. A similar argument for $i = k$ shows $x^*(\partial(L_k; L_{k+1})) \geq 1 - \tau$. So it remains to consider $1 < i < k$. Define the following quantities, some of which can be zero.

- $A = x^*(\partial(L_i; L_{i+1}))$
- $B = x^*(\partial(L_i; L_{\geq i+2}))$
- $C = x^*(\partial(L_{\leq i-1}; L_{i+1}))$

We have

$$1 \leq x^*(\partial^+(L_i)) = A + B + x^*(\partial(L_i; L_{\leq i-1})) \leq A + B + \tau,$$

since $\partial(L_i; L_{\leq i-1}) \subseteq \partial^-(U_{i-1})$ and U_{i-1} is τ -narrow. Similarly

$$1 \leq x^*(\partial^-(L_{i+1})) = A + C + x^*(\partial(L_{\geq i+2}; L_{i+1})) \leq A + C + \tau.$$

Summing these two inequalities yields $2 \leq A + (A + B + C) + 2\tau \leq A + (1 + \tau) + 2\tau$ where we have used $A + B + C \leq x^*(\partial^+(U_i)) \leq 1 + \tau$. Rearranging shows $A \geq 1 - 3\tau$. \square

Now, recall that $\kappa(x^*)$ denotes the assignment of arc weights to the graph $H = (V, E)$ from the previous section obtained by “removing” the directions from arcs in A . We prove that the restriction of $\kappa(x^*)$ to any L_i almost satisfies the partition inequalities that characterize the convex hull of connected graphs. For a partition $\pi = \{W_1, \dots, W_\ell\}$, we let $\partial(\pi)$ denote the set of edges whose endpoints lie in two different sets in the partition.

Lemma 2.5. *For any $1 \leq i \leq k + 1$ and any partition $\pi = \{W_1, \dots, W_\ell\}$ of L_i , we have $\kappa(x^*)(\partial(\pi)) \geq \ell - 1 - 2\tau$.*

Proof. Since $L_1 = \{s\}$ and $L_{k+1} = \{t\}$, then there is nothing to prove for $i = 1$ or $i = k + 1$. So, we suppose $1 < i < k + 1$.

Consider the quantity $X = \sum_{j=1}^{\ell} x^*(\partial^+(W_j)) + x^*(\partial^-(W_j))$. On one hand, since neither s nor t are in any W_j , then $x^*(\partial^+(W_j)) = x^*(\partial^-(W_j)) \geq 1$ so $X \geq 2\ell$. On the other hand, X counts each arc between two partitions in π exactly twice and each arc with one end in L_i and the other not in L_i precisely once. So, $X = 2\kappa(x^*)(\partial(\pi)) + x^*(\partial^+(L_i)) + x^*(\partial^-(L_i))$.

Notice that $\partial^+(L_i)$ and $\partial^-(L_i)$ are disjoint subsets of $\partial^+(U_{i-1}) \cup \partial^-(U_{i-1}) \cup \partial^+(U_i) \cup \partial^-(U_i)$. So, since both U_{i-1} and U_i are τ -narrow, then $x(\partial^+(L_i)) + x(\partial^-(L_i)) < 2 + 4\tau$. This shows $2\ell \leq X < 2\kappa(x^*)(\partial(\pi)) + 2 + 4\tau$ which, after rearranging, is what we wanted to show. \square

The following corollary will be useful.

Corollary 2.6. *For any partition π of L_i , we have $\frac{\kappa(x^*)(\partial(\pi))}{1-2\tau} \geq |\pi| - 1$.*

Proof. From Lemma 2.5, we have $\frac{\kappa(x^*)(\partial(\pi))}{1-2\tau} \geq \frac{|\pi| - 1 - 2\tau}{1-2\tau} \geq |\pi| - 1$ for any $|\pi| \geq 2$. \square

Finally, to efficiently implement the arguments in the proof of Lemma 2.2, we need to be able to efficiently find all τ -narrow cuts U_i . This is done by a standard recursive algorithm that exploits the fact that the cuts are nested.

Lemma 2.7. *There is a polynomial-time algorithm to find all τ -narrow $s - t$ cuts.*

Proof. Consider following standard recursive algorithm. As input, the routine is given a directed graph $H = (V', A')$ with arc weights x_a^* and distinct nodes s', t' where both $\{s'\}$ and $V' \setminus \{t'\}$ are τ -narrow. Say a τ -narrow cut U in H is non-trivial if $U \neq \{s'\}$ and $U \neq V' \setminus \{t'\}$. The claim is that the procedure will find all non-trivial τ -narrow $s - t$ cuts of H , provided that they are nested.

The procedure works as follows. If there are non-trivial τ -narrow $s - t$ cuts in H , then there are nodes $u, v \in V' \setminus \{s', t'\}$ such that some τ -narrow $s' - t'$ cut U has $\{s', u\} \subseteq U \subseteq V' \setminus \{t', v\}$. So, the procedure tries all $O(|V'|^2)$ pairs of distinct nodes u, v , contracts both $\{s', u\}$ and $\{t', v\}$ to a single node and determines if the minimum cut separating these contracted nodes has capacity less than $1 + \tau$. If such a cut U was found for some u, v , the algorithm makes two recursive calls, one with the contracted graph $H[V'/U]$ with start node being the contraction of U and end node being t' , and the other with the contracted graph $H[V'/(V' \setminus U)]$ with start node s' and end node being the contraction of $V' \setminus U$. After both recursive calls complete, the algorithm returns all τ -narrow cuts found by these two recursive calls (of course, after expanding the contracted nodes) and the τ -narrow cut U itself. If such a cut U was not found over all choices of u, v , then the algorithm returns nothing since there are no non-trivial τ -narrow $s' - t'$ cuts in H .

It is easy to see that a non-trivial τ -narrow cut in either contracted graph corresponds to a τ -narrow cut in H . On the other hand, if the τ -narrow $s' - t'$ cuts are nested in H , then every non-trivial τ -narrow $s' - t'$ cut apart from U itself corresponds to a non-trivial τ -narrow cut in exactly one of $H[V'/U]$ or $H[V'/(V' \setminus U)]$. Also, the τ -narrow cuts in both contracted graphs remain nested. So, the recursive procedure finds all non-trivial τ -narrow cuts of H . The number of recursive calls is at most the number of non-trivial τ -narrow cuts, and this is at most $|V'|$ because the cuts are nested so it is an efficient algorithm. We call this algorithm initially with graph G , start node s and end node t . [Lemma 2.3](#) implies the τ -narrow $s - t$ cuts of G are nested so the recursive procedure finds all non-trivial τ -narrow cuts of G . Adding these to $\{s\}$ and $V \setminus \{t\}$ gives all τ -narrow cuts of G . \square

We are now in a position to prove [Lemma 2.2](#), the main result of this section.

Proof of Lemma 2.2. The claimed vector z can be described by linear constraints: indeed, consider the following LP on the variables z where constraints (5) imply that $\kappa(z)$ is in the convex hull of spanning connected graphs [[16](#), Corollary 50.8a].¹

$$\kappa(z)(\partial(\pi)) \geq |\pi| - 1 \quad \forall \text{ partitions } \pi \text{ of } V \quad (5)$$

$$z_a \leq \frac{1}{1-3\tau} x_a^* \quad \forall a \in A \quad (6)$$

$$z(\partial^+(U_i)) = 1 \quad \forall \tau\text{-narrow } s\text{-}t \text{ cuts } U_i \quad (7)$$

$$z(\partial^-(U_i)) = 0 \quad \forall \tau\text{-narrow } s\text{-}t \text{ cuts } U_i \quad (8)$$

$$z_a \geq 0 \quad \forall a \in A \quad (9)$$

We demonstrate a feasible z as follows.

$$z_a = \begin{cases} \frac{x_a^*}{x^*(\partial(L_i; L_{i+1}))} & \text{if } a \in \partial(L_i; L_{i+1}) \text{ for some } i; \\ \frac{x_a^*}{1-2\tau} & \text{if } a \in E[L_i] \text{ for some } i; \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

We claim that this solution z satisfies the above constraints. Constraints (8) and (9) are satisfied by construction. Constraint (6) follows from [Lemma 2.4](#) for edges in $\partial(L_i; L_{i+1})$ and by construction for rest of the edges. For constraint (7), note that

$$z(\partial^+(U_i)) = z(\partial(L_i; L_{i+1})) + z(\partial^+(U_i) \setminus \partial(L_i; L_{i+1})) = \frac{x^*(\partial(L_i; L_{i+1}))}{x^*(\partial(L_i; L_{i+1}))} + 0 = 1.$$

To complete the proof, we now show constraints (5) holds. It suffices to show that $\kappa(z)$ can be decomposed as a convex combination of characteristic vectors of connected graphs. For $1 \leq i \leq k+1$, let z^i denote the restriction of $\kappa(z)$ to edges whose endpoints are both contained in L_i . Then [Corollary 2.6](#), constraints (9), and [[16](#), Corollary 50.8a] imply that z^i can be decomposed as a convex combination of integral vectors, each of which corresponds to an edge set that is connected on L_i . Next, let z' denote the restriction of $\kappa(z)$ to edges whose endpoints are both contained in some common L_i . Since the sets $E(L_1), \dots, E(L_{k+1})$ are disjoint, we have that $z' = \sum_i z^i$ (where the addition is component-wise). Furthermore, z' , being the sum of

¹The statement of [Lemma 2.2](#) makes a claim about $\kappa(z)$ being in the convex hull of spanning *trees* and not spanning *connected graphs*. However, the equivalent statement for spanning trees will follow by dropping some edges from the connected subgraphs in the decomposition of z to get spanning trees. Constraints (7) and (8) will still be satisfied since we retain connectivity.

the z^i vectors, can be decomposed as a convex combination of integral vectors corresponding to edge sets E' such that the connected components of the graph $H' = (V, E')$ are precisely the sets $\{L_i\}_{i=1}^{k+1}$.

Next, let z'' denote the restriction of $\kappa(z)$ to edges contained in one such $\partial(L_i; L_{i+1})$. We also note that the sets $\partial(L_1; L_2), \dots, \partial(L_k; L_{k+1})$ are disjoint. By construction, we have $z''(\partial(L_i; L_{i+1})) = 1$ for each $1 \leq i \leq k$ so we may decompose z'' as a convex-combination of integral vectors, each of which includes precisely one edge across each $\partial(L_i; L_{i+1})$.

Now, adding any integral point y' in the decomposition of z' to any integral point y'' in the decomposition of z'' results in an integral vector that corresponds to a connected graph: each L_i is connected by y' and consecutive L_i are connected by y'' . By construction of z , we have $\kappa(z) = z' + z''$ so we may write z as a convex combination of characteristic vectors of connected graphs, each of which satisfies constraints (5).

To see why z can be found efficiently, we first compute all τ -narrow cuts using [Lemma 2.7](#). Then z is easy to compute in equation (10). Finally, [16, Corollary 51.6a] implies the decomposition of $\kappa(z)$ into a convex combination of connected graphs can be done efficiently, so the arguments in the footnote to reduce z such that $\kappa(z)$ is in the spanning tree polytope can also be done efficiently. \square

3 Obtaining an s - t Path

Having transformed the optimal LP solution x^* into the new vector z (as in [Lemma 2.2](#)) without increasing it too much in any coordinate, we now sample a random tree such that it has a small total cost, and that the tree does not cross any cut much more than prescribed by x^* . Finally we add some arcs to this tree (without increasing its cost much) so that it is Eulerian at all nodes except $\{s, t\}$, and hence gives us an Eulerian s - t walk. By the triangle inequality, shortcutting this walk past repeated nodes yields a Hamiltonian s - t path of no greater cost. While this general approach is similar to that used in [2], some new ideas are required because we are working with the LP for ATSP—in particular, only one unit of flow is guaranteed to cross s - t cuts, which is why we needed to deal with narrow cuts in the first place. The details appear in the rest of this section.

3.1 Sampling a Tree

For a collection of arcs $\mathcal{A} \subset A$, we say \mathcal{A} is α -thin with respect to x^* if $|\mathcal{A} \cap \partial^+(U)| \leq \alpha x^*(\partial^+(U))$ for every $\emptyset \subsetneq U \subsetneq V$. The set \mathcal{A} is also β -approximate with respect to x^* if the total cost of all arcs in \mathcal{A} is at most β times the cost of x^* —i.e., $\sum_{a \in \mathcal{A}} c_a \leq \beta \sum_{a \in A} c_a x_a^*$. The reason we are deviating from the undirected to the directed setting is that the orientation of the arcs across each τ -narrow cut will be important when we sample a random “tree”.

Lemma 3.1. *Let $\tau \in [0, 1/4]$. Let $\beta = \frac{3}{1-3\tau}$ and $\alpha = (2 + \frac{1}{\tau}) \cdot \frac{24 \log n}{\log \log n}$. For sufficiently large n , there is a randomized, polynomial time algorithm that, with probability at least $1/2$, finds an α -thin and β -approximate (with respect to x^*) collection of arcs \mathcal{A} that is weakly connected and satisfies $|\mathcal{A} \cap (\partial^+(U))| = 1$ and $|\mathcal{A} \cap (\partial^-(U))| = 0$ for each τ -narrow s - t cut U .*

Proof. Let z be a vector as promised by [Lemma 2.2](#). From $\kappa(z)$, randomly sample a set of arcs \mathcal{A} whose undirected version \mathcal{T} is a spanning tree on V . This should be done from any distribution with the following two properties:

- (i) (*Correct Marginals*) $\Pr[e \in \mathcal{T}] = \kappa(z)_e$
- (ii) (*Negative Correlation*) For any subset of edges $F \subseteq E$, $\Pr[F \subseteq \mathcal{T}] \leq \prod_{e \in F} \Pr[e \in \mathcal{T}]$

This can be obtained using, for example, the swap rounding approach for the spanning tree polytope given by Chekuri et al. [7]. As in [2], the negative correlation property implies the following theorem. The proof is found in Section 4.

Theorem 3.2. *The tree \mathcal{T} is α -thin with high probability.*

By Lemma 2.2(b), property (i) of the random sampling, and Markov's inequality, we get that \mathcal{A} (from Lemma 3.1) is $\frac{3}{1-3\tau}$ -approximate with respect to x^* with probability at least $2/3$. By a trivial union bound, for large enough n we have with probability at least $1/2$ that \mathcal{A} is both α -thin and β -approximate with respect to x^* . It is also weakly connected—i.e., the undirected version of \mathcal{A} (namely, \mathcal{T}) connects all vertices in V .

The statement for τ -narrow s - t cuts follows from the fact that z satisfies Lemma 2.2(c). That is, \mathcal{A} contains no arcs of $\partial^-(U)$, since $z(\partial^-(U)) = 0$ (for U being a τ -narrow s - t cut). But since \mathcal{T} is a spanning tree, \mathcal{A} must contain at least one arc from $\partial^+(U)$. Finally, since $z(\partial^+(U))$ is exactly 1, then any set of arcs supported by this distribution we use must have precisely one arc from $\partial^+(U)$. \square

3.2 Augmenting to an Eulerian s - t Walk

Finally, we wrap up by augmenting the set of arcs \mathcal{A} to an Eulerian s - t walk. For this, we use Hoffman's circulation theorem, as in [2], which we recall here for convenience (see, e.g. [16, Theorem 11.2]):

Theorem 3.3. *Given a directed flow network $D = (V, A)$, with each arc having a lower bound ℓ_a and an upper bound u_a (and $0 \leq \ell_a \leq u_a$), there exists a circulation $f : A \rightarrow \mathbb{R}_+$ satisfying $\ell_a \leq f(a) \leq u_a$ for all arcs a if and only if $\ell(\partial^+(U)) \leq u(\partial^-(U))$ for all $U \subseteq V$. Moreover, if the ℓ and u are integral, then the circulation f can be taken integral.*

Set lower bounds $\ell : A \rightarrow \{0, 1\}$ on the arcs by:

$$\ell_a = \begin{cases} 1 & \text{if } a \in \mathcal{A} \text{ or } a = ts \\ 0 & \text{otherwise} \end{cases}$$

For now, we set an upper bound of 1 on arc ts and leave all other arc upper bounds at ∞ . We compute the minimum cost circulation satisfying these bounds (we will soon see why one must exist). Since the bounds are integral and since \mathcal{A} is weakly connected, this circulation gives us a directed Eulerian graph. Furthermore, since $u_{ts} = \ell_{ts} = 1$, the ts arc must appear exactly once in this Eulerian graph. Our final Hamiltonian s - t path is obtained by following an Eulerian circuit, removing the single ts arc from this circuit to get an Eulerian s - t walk, and finally shortcutting this walk past repeated nodes. The cost of this Hamiltonian path will be, by the triangle inequality, at most the cost of the circulation minus the cost of the ts arc.

Finally, we need to bound the cost of the circulation (and also to prove one exists). To this end, we will impose further upper bounds $u : A \rightarrow \mathbb{R}_{\geq 0}$ as follows:

$$u_a = \begin{cases} 1 & \text{if } a = ts \\ 1 + (1 + \tau^{-1})\alpha x_a^* & \text{if } a \in \mathcal{A} \\ (1 + \tau^{-1})\alpha x_a^* & \text{otherwise} \end{cases}$$

We use Hoffman's circulation theorem to show that a circulation f exists satisfying these bounds ℓ and u (The calculations appear in the next paragraph.) Since u is no longer integral, the circulation f might

not be integral, but it does demonstrate that a circulation exists where each arc $a \neq ts$ is assigned at most $(1 + \tau^{-1})\alpha x^* a$ more flow in the circulation than the number of times it appears in \mathcal{A} . Consequently, it shows that the minimum cost circulation g in the setting where we only had a non-trivial upper bound of 1 on the arc ts can be no more expensive (since there are fewer constraints), and that circulation g can be chosen to be integral. The cost of circulation g is at most the cost of f , which is at most

$$\sum_{a \in A} c_a u_a = \sum_{a \in \mathcal{A}} c_a + (1 + \tau^{-1})\alpha \sum_{a \in A} c_a x_a^* + c_{ts}.$$

Subtracting the cost of the ts arc (since we drop it to get the Hamilton path) and recalling that \mathcal{A} is $\frac{3}{1-3\tau}$ -approximate with respect to x^* (and hence $\sum_{a \in \mathcal{A}} c_a \leq \frac{3}{1-3\tau} \sum_{a \in A} c_a x_a^*$, we get that the final Hamiltonian path has cost at most

$$\left(\frac{3}{1-3\tau} + (1 + \tau^{-1})\alpha \right) \sum_{a \in A} c_a x_a^*,$$

and hence $O\left(\frac{\log n}{\log \log n}\right)$ times the cost of the LP relaxation for $\tau = 1/4$. This proves the claim that the cost of the s - t path we found is $O\left(\frac{\log n}{\log \log n}\right)$ times the LP value, with constant probability, and completes the proof of [Theorem 1.1](#).

One detail remains: we need to verify the conditions of [Theorem 3.3](#) for the bounds ℓ and u . Firstly, it is clear by definition that $\ell_a \leq u_a$ for each arc a . Now we need to check $\ell(\partial^+(U)) \leq u(\partial^-(U))$ for each cut U . This is broken into four cases (where saying U is a u - v cut means $u \in U, v \notin U$).

1. U is a τ -narrow s - t cut. Then $\ell(\partial^+(U)) = 1$, since \mathcal{A} contains only one arc in $\partial^+(U)$. But $1 = u_{ts} \leq u(\partial^-(U))$.
2. U is an s - t cut, but not τ -narrow. Then by the α -thinness of \mathcal{A} ,

$$\ell(\partial^+(U)) \leq \alpha x^*(\partial^+(U)) = \alpha x^*(\partial^-(U)) + \alpha.$$

On the other hand,

$$u(\partial^-(U)) \geq (1 + \tau^{-1})\alpha x^*(\partial^-(U)) = \alpha x^*(\partial^-(U)) + \tau^{-1}\alpha x^*(\partial^-(U)) \geq \alpha x^*(\partial^-(U)) + \alpha$$

where the last inequality used the fact that $x^*(\partial^-(U)) \geq \tau$.

3. U is a t - s cut. Then

$$\ell(\partial^+(U)) \leq 1 + \alpha x^*(\partial^+(U)) = 1 + \alpha x^*(\partial^-(U)) - \alpha \leq \alpha x^*(\partial^-(U)),$$

the last inequality using that $\alpha \geq 1$. Moreover

$$u(\partial^-(U)) \geq (1 + \tau^{-1})\alpha x^*(\partial^-(U)) \geq \alpha x^*(\partial^-(U)).$$

Then $\ell(\partial^+(U)) \leq u(\partial^-(U))$.

4. U does not separate s from t . Then

$$\ell(\partial^+(U)) \leq \alpha x^*(\partial^+(U)) = \alpha x^*(\partial^-(U)) \leq (1 + \tau^{-1})\alpha x^*(\partial^-(U)) \leq u(\partial^-(U))$$

4 Guaranteeing α -Thinness

We prove [Theorem 3.2](#) in this section. Recall that α -thin means the number of arcs chosen from $\partial^+(U)$ should not exceed $\alpha x^*(\partial^+(U))$ (so a directed version). Let $\alpha := \left(2 + \frac{1}{\tau}\right) \cdot \frac{24 \log n}{\log \log n}$ where the logarithm is the natural logarithm. Recall that \mathcal{A} is the set of arcs found with corresponding undirected spanning tree \mathcal{T} . By the first property of the distribution (preservation of marginals on singletons) we have for each $\emptyset \subsetneq U \subsetneq V$ that $\mathbb{E}[|\partial_{\mathcal{T}}(U)|] = \kappa(z)(\partial(U))$.²

Since we have negative correlation on subsets of items, then we can apply a Chernoff bound. For notational simplicity, let $z' := \kappa(z)$. Then we use

$$\Pr[|\partial_{\mathcal{T}}(U)| \geq (1 + \delta)z'(\partial(U))] \leq \left(\frac{e^\delta}{(1 + \delta)^{(1 + \delta)}} \right)^{z'(\partial(U))}.$$

Let $\beta := \frac{6 \log n}{\log \log n}$ (again using the natural logarithm) and let $\delta := \beta - 1$. For large enough n , the above expression is bounded (in a manner similar to [2]) by

$$\left(\frac{e}{\beta} \right)^{\beta z'(\partial(U))} \leq e^{-z'(\partial(U))5 \log n} = n^{-5z'(\partial(U))}.$$

However, for any graph, there are at most n^{2l} cuts whose capacity is at most l times the capacity of the minimum cut [12]. Since the minimum cut with capacities z' is 1, then there are at most n^{2l} cuts of the undirected graph H with capacity (under z') at most l . Another way to view this is that there are at most $n^{2(l+1)}$ cuts whose capacity is between l and $l + 1$. For each such cut U , the previous analysis shows that probability that $|\partial_{\mathcal{T}}(U)| > (1 + \delta)z'(\partial(U))$ is at most n^{-5l} . Thus, by the union bound, the probability that $|\partial_{\mathcal{T}}(U)| > (1 + \delta)z'(\partial(U))$ for some $\emptyset \subsetneq U \subsetneq V$ is bounded by

$$\sum_{i=1}^{\infty} n^{2(i+1)} \cdot n^{-5i} \leq \sum_{i=1}^{\infty} n^{-i} = \frac{1}{n-1}$$

Since $|\partial_{\mathcal{A}}^+(U)| \leq |\partial_{\mathcal{T}}(U)|$, then we have just seen that with probability at least $1 - \frac{1}{n-1}$ that there is no $\emptyset \subsetneq U \subsetneq V$ with $|\partial_{\mathcal{A}}^+(U)| > \beta z'(\partial(U))$. Now, there are three additional cases to consider:

- $x^*(\partial^+(U)) < \tau$ (so $t \in U$ and $s \notin U$). We actually ignore the above analysis in this case and simply use the fact that $z(\partial^+(U)) = 0$ is guaranteed by construction of z . Then we trivially have $|\partial_{\mathcal{A}}^+(U)| = 0 \leq \alpha x^*(\partial^+(U))$
- $t \in U$ and $s \notin U$, but $x^*(\partial^+(U)) \geq \tau$. Since $\tau \leq 1/4$, then $\frac{1}{1-3\tau} \leq 4$ so $z \leq 4x^*$. We have

$$\begin{aligned} |\partial_{\mathcal{A}}^+(U)| &\leq \beta z'(\partial(U)) \\ &\leq 4\beta(x^*(\partial^+(U)) + x^*(\partial^-(U))) \\ &= 4\beta(2x^*(\partial^+(U)) + 1) \\ &= 8\beta x^*(\partial^+(U)) + 4\beta \\ &\leq 8\beta x^*(\partial^+(U)) + \frac{4\beta}{\tau} x^*(\partial^+(U)) \\ &= \alpha x^*(\partial^+(U)) \end{aligned}$$

²Here we use $\partial_{\mathcal{T}}(U)$ to denote the set $\partial(U) \cap \mathcal{T}$; similarly we will let $\partial_{\mathcal{A}}^+(U)$ denote $\partial^+(U) \cap \mathcal{A}$, etc.

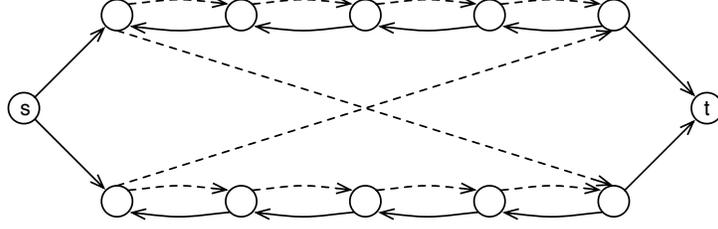


Figure 1: The graph G_r with $r = 5$. The solid edges have cost 1 and the dashed edges have cost 0.

- Finally, if it is not the case that $\{t\} \subseteq U \subseteq V \setminus \{s\}$, then $x^*(\partial^+(U)) + x^*(\partial^-(U)) \leq 2x^*(\partial^+(U))$ so we get (using $z \leq 4x^*$ again)

$$\begin{aligned}
|\partial_{\mathcal{A}}^+(U)| &\leq \beta z'(\partial(U)) \\
&\leq 4\beta(x^*(\partial^+(U)) + x^*(\partial^-(U))) \\
&\leq 8\beta x^*(\partial^+(U)) \\
&\leq \alpha x^*(\partial^+(U))
\end{aligned}$$

Summarizing, for sufficiently large n we have with probability at least $1 - \frac{1}{n-1}$ that

$$|\partial_{\mathcal{A}}^+(U)| \leq \alpha x^*(\partial^+(U)) = \Theta\left(\frac{\log n}{\log \log n}\right) x^*(\partial^+(U)).$$

That is, \mathcal{A} is α -thin with high probability.

5 A Simple Integrality Gap Example

In this section, we show that the integrality gap of the subtour elimination LP (*ATSP*) is at least 2. This result can also be inferred from the integrality gap of 2 for the ATSP tour problem [5], but our construction is relatively simpler.

For a fixed integer $r \geq 1$, consider the directed graph G_r defined below (and illustrated in Figure 1). The vertices of G_r are $\{s, t\} \cup \{u_1, \dots, u_r\} \cup \{v_1, \dots, v_r\}$; the edges are as follows:

- $\{su_1, sv_1, u_r t, v_r t\}$, each with cost 1,
- $\{u_1 v_r, v_1 u_r\}$, each with cost 0,
- $\{u_{i+1} u_i \mid 1 \leq i < r\} \cup \{v_{i+1} v_i \mid 1 \leq i < r\}$, each with cost 1,
- and $\{u_i u_{i+1} \mid 1 \leq i < r\} \cup \{v_i v_{i+1} \mid 1 \leq i < r\}$, each with cost 0.

Let F_r denote the ATSP instance obtained from the metric completion of G_r .

Lemma 5.1. *The integrality gap of the LP ATSP on the instance F_r is at least $2 - o(1)$.*

Proof. It is easy to verify that assigning $x_a = 1/2$ to each arc that originally appeared in G_r is a valid LP solution. Indeed, the degree constraints are immediate, and there are two edge-disjoint paths from s to every other node in G_r (so there must be at least 2 arcs exiting any subset containing s) so the cut constraints are also satisfied. The total cost of this LP solution is $r + 1$.

On the other hand, we claim that the cost of any Hamiltonian s - t path in L_r , which corresponds to a spanning s - t walk W in G_r , is at least $2r - 1$. This shows an integrality gap of $\frac{2r-1}{r+1} = 2 - o(1)$.

To lower-bound the length of any spanning s - t walk, we first argue that the walk W can avoid using at most one of the unit cost edges of the form $u_{i+1}u_i$ or $v_{i+1}v_i$. Indeed, any u_r - v_r walk must use edges $u_{i+1}u_i$ for every $1 \leq i < r$. Similarly, every v_r - u_r walk must use all edges of the form $v_{i+1}v_i$. One of u_r and v_r is visited before the other, so either all of the $u_{i+1}u_i$ edges or all of the $v_{i+1}v_i$ edges are used by W . Now suppose, without loss of generality, that W does not use the edges $u_{i+1}u_i$ and $u_{j+1}u_j$ for $1 \leq i < j < r$. Every u_{i+1} - v_r walk uses edge $u_{i+1}u_i$ and every v_r - u_{i+1} walk uses edge $u_{j+1}u_j$. Since one of u_{i+1} or v_r must be visited by W before the other, then W cannot avoid both $u_{i+1}u_i$ and $u_{j+1}u_j$ which contradicts our assumption.

Thus, W must use all but at most one of the $2r - 2$ unit cost edges in $\{u_{i+1}u_i \mid 1 \leq i < r\} \cup \{v_{i+1}v_i \mid 1 \leq i < r\}$. Moreover, W must also use one of the arcs exiting s and one of the arcs entering t , so the cost of W is at least $2r - 1$. (In fact, the walk

$$\langle s, u_1, v_r, v_{r-1}, \dots, v_1, u_r, u_{r-1}, \dots, u_3, u_2, u_3, \dots, u_r, t \rangle$$

is of length exactly $2r - 1$, so this argument is tight.) □

6 Conclusion

In this paper we showed that the integrality gap for the ATSP problem is $O(\frac{\log n}{\log \log n})$. In fact, our proof also shows an integrality gap of α for ATSP whenever we can construct a procedure which takes a point $y \in \mathbb{R}^{|E|}$ in the spanning tree polytope of an undirected (multi)graph $H = (V, E)$ and outputs a tree T that is (a) α -thin, and (b) also satisfies $|T \cap \partial(U)| = 1$ for any cut U where $y(\partial(U)) = 1$. We also showed a simpler construction achieving a lower bound of 2 for the subtour elimination LP.

Acknowledgments.

We thank V. Nagarajan for enlightening discussions in the early stages of this project. Z.F. and A.G. also thank A. Vetta and M. Singh for their generous hospitality. Part of this work was done when A.G. was visiting the IEOR Department at Columbia University. Part of this work was done when M.S. was at McGill University.

References

- [1] H.-C. An, R. D. Kleinberg, and D. B. Shmoys. Improving Christofides' algorithm for the s - t path TSP. In *Proc. 44th ACM Symp. on Theory of Computing*, 2012.
- [2] A. Asadpour, M. X. Goemans, A. Mađry, S. Oveis Gharan, and A. Saberi. An $O(\log n / \log \log n)$ -approximation algorithm for the asymmetric traveling salesman problem. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 379–389, Philadelphia, PA, 2010. SIAM.
- [3] M. Bateni and J. Chuzhoy. Approximation algorithms for the directed k -tour and k -stroll problems. In *APPROX-RANDOM*, pages 25–38, 2010.
- [4] M. Bläser. A new approximation algorithm for the asymmetric TSP with triangle inequality. *ACM Trans. Algorithms*, 4(4):Art. 47, 15, 2008.
- [5] M. Charikar, M. X. Goemans, and H. Karloff. On the integrality ratio for the asymmetric traveling salesman problem. *Math. Oper. Res.*, 31(2):245–252, 2006.

- [6] C. Chekuri and M. Pál. An $O(\log n)$ approximation ratio for the asymmetric traveling salesman path problem. *Theory Comput.*, 3:197–209, 2007.
- [7] C. Chekuri, J. Vondrák, and R. Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *FOCS*, pages 575–584, 2010.
- [8] U. Feige and M. Singh. Improved approximation ratios for traveling salesperson tours and paths in directed graphs. In *APPROX-RANDOM*, pages 104–118, 2007.
- [9] A. M. Frieze, G. Galbiati, and F. Maffioli. On the worst-case performance of some algorithms for the asymmetric traveling salesman problem. *Networks*, 12(1):23–39, 1982.
- [10] Z. Friggstad, M. R. Salavatipour, and Z. Svitkina. Asymmetric traveling salesman path and directed latency problems. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 419–428, Philadelphia, PA, 2010. SIAM.
- [11] H. Kaplan, M. Lewenstein, N. Shafir, and M. Sviridenko. Approximation algorithms for asymmetric TSP by decomposing directed regular multigraphs. *J. ACM*, 52(4):602–626, 2005.
- [12] D. R. Karger and C. Stein. A new approach to the minimum cut problem. *J. ACM*, 43(4):601–640, 1996.
- [13] F. Lam and A. Newman. Traveling salesman path problems. *Math. Program.*, 113(1, Ser. A):39–59, 2008.
- [14] V. Nagarajan and R. Ravi. Poly-logarithmic approximation algorithms for directed vehicle routing problems. In *APPROX-RANDOM*, pages 257–270, 2007.
- [15] V. Nagarajan and R. Ravi. The directed minimum latency problem. In *Approximation, randomization and combinatorial optimization*, volume 5171 of *Lecture Notes in Comput. Sci.*, pages 193–206. Springer, Berlin, 2008.
- [16] A. Schrijver. *Combinatorial optimization. Polyhedra and efficiency.*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003.
- [17] A. Sebö and J. Vygen. Shorter tours by nicer ears: $7/5$ -approximation for graphic TSP, $3/2$ for the path version, and $4/3$ for two-edge-connected subgraphs. *CoRR*, abs/1201.1870, 2012.