

A Constant-Factor Approximation for Directed Latency in Quasi-Polynomial Time

Zachary Friggstad

Department of Computer Science, University of Alberta, Edmonton, Canada
zacharyf@ualberta.ca

Chaitanya Swamy

Department of Combinatorics and Optimization, University of Waterloo, Canada
cswamy@uwaterloo.ca

Abstract

We consider the *directed minimum latency problem* (DirLat), wherein we seek a path P visiting all points (or clients) in a given asymmetric metric starting at a given root node r , so as to minimize the sum of the client waiting times, where the waiting time of a client v is the length of the r - v portion of P . We give the *first constant-factor approximation* guarantee for DirLat, but in quasi-polynomial time. Previously, a polynomial-time $O(\log n)$ -approximation was known [12], and no better approximation guarantees were known even in quasi-polynomial time.

A key ingredient of our result, and our chief technical contribution, is an extension of a recent result of [17] showing that the integrality gap of the natural Held-Karp relaxation for *asymmetric TSP-Path* (ATSP_P) is at most a constant, which itself builds on the breakthrough similar result established for *asymmetric TSP* (ATSP) by Svensson et al. [25]. We show that the integrality gap of the Held-Karp relaxation for ATSP_P is bounded by a constant even if the cut requirements of the LP relaxation are relaxed from $x(\delta^{\text{in}}(S)) \geq 1$ to $x(\delta^{\text{in}}(S)) \geq \rho$ for some constant $1/2 < \rho \leq 1$.

We also give a better approximation guarantee for the *minimum total-regret problem*, where the goal is to find a path P that minimizes the total time that nodes spend in excess of their shortest-path distances from r , which can be cast as a special case of DirLat involving so-called regret metrics.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis

Keywords and phrases Approximation Algorithms, Directed Latency, TSP

Digital Object Identifier 10.4230/LIPIcs.ESA.2020.52

Related Version A full version of the paper is available at <https://arxiv.org/abs/1912.06198>.

Funding Zachary Friggstad: Supported by the Canada Research Chairs program and an NSERC Discovery grant.

Chaitanya Swamy: Supported in part by NSERC grant 327620-09 and an NSERC Discovery Accelerator Supplement Award.

1 Introduction

Vehicle-routing problems form a rich class of combinatorial-optimization problems that find applications in a wide variety of settings, and have been extensively studied in the Operations Research and Computer Science communities (see, e.g., [26]). These problems typically involve designing routes for vehicles to service a given underlying set of clients in the most time- and/or cost-effective fashion. A fundamental problem in this genre is the *minimum latency problem* (MLP), also known as the *traveling repairman problem* or the *delivery man problem* [1, 19, 10, 5], wherein, adopting a client-oriented perspective, we seek a route starting at a given root node and visiting all client nodes that minimizes the sum of client waiting times (or equivalently, the average client waiting time).¹

¹ In contrast, the path-version of TSP can be seen as minimizing the *maximum* client waiting time.



We investigate *directed* MLP (DirLat), i.e., MLP in directed (or asymmetric) metrics. Formally, we are given an asymmetric metric space $(V \cup \{r\}, c)$, where V is a set of client nodes, r is a root or depot node, and $c = \{c_{u,v}\}_{u,v \in V \cup \{r\}}$ specifies the asymmetric metric: in particular, for any $u, v, w \in V \cup \{r\}$, we have $c_{u,u} = 0$, $c_{u,v} \geq 0$, and $c_{u,v} \leq c_{u,w} + c_{w,v}$. The goal is to find a Hamiltonian path P starting at the depot r to minimize $\sum_{v \in V} c_P(v)$, where $c_P(v)$ is the cost of the $r \rightsquigarrow v$ subpath of P and is interpreted as the waiting time or *latency* of node v . Throughout, we let n denote $|V|$.

Whereas we have a reasonably good understanding of MLP in undirected (i.e., symmetric) metrics – a constant-factor approximation is known (see [9] and the references therein) and recent work has also led to LP-based approaches [7, 22] for the problem – there are significant gaps in our understanding of directed MLP: the approximation factor has remained stagnant at $O(\log n)$ [12] for close to a decade, and it is not known if $\log n$ (or any super-constant function of n) constitutes a real inapproximability barrier for the problem.

2 Our contributions

Our main contribution is to provide the *first constant-factor approximation* guarantee for DirLat, albeit in quasi-polynomial time, i.e., $O(n^{O(\log n)})$ time. This provides the first concrete indication that $\log n$ is unlikely to be an inapproximability barrier for DirLat (unless $NP \subseteq DTIME [n^{O(\log n)}]$).

► **Theorem 1.** *There is an $O(1)$ -approximation for DirLat running in $O(n^{O(\log n)})$ time.*

Our algorithm is based on a natural time-indexed linear programming (LP) relaxation that is similar to, and inspired by, the approach taken in [22] for undirected MLP. Roughly speaking, our LP (LP-Lat) utilizes variables for (v, t) pairs where $v \in V$ is a node to be visited and t is the time they should be visited, and other variables indicating the edges present on various prefixes of the optimal path. As is typical for minimum-latency problems, we utilize the LP to find rooted paths of geometrically increasing lengths and stitch them together. However, with asymmetric metrics, *both* steps present significant challenges. A key technical contribution underlying our result is a procedure for achieving the former step, namely a way of rounding the fractional prefix of the optimal path of length t to obtain a rooted path of length $O(t)$. We achieve this by generalizing some recent work by [17] on the integrality gap for *asymmetric s-t TSP-path* (ATSP). We show that the integrality gap of a weakening of the standard LP-relaxation, where we require non- s - t cuts to only be covered to some extent strictly larger than 1 still remains a constant (see Theorem 2 below).

An interesting special case of DirLat, involves the notion of *regret* of a client: the regret of a client v lying on a rooted path P is defined as $c_P(v) - c_{r,v}$; that is, regret measures the time that node v spends waiting *in excess* of its its least possible waiting time. The notion of regret can be seen as a nuanced and better way of measuring the (dis)satisfaction of a client than the standard measure of simply considering the waiting time of a client. The latter does not differentiate between clients located at different distances from the depot and their varying expectations, and fails to take into account that a client closer to the depot that incurs a larger delay than further-away clients may face a greater level of dissatisfaction. A natural problem that arises is to find a path that minimizes the total regret of clients, i.e., to minimize $\sum_{v \in V} (c_P(v) - c_{r,v})$ (or equivalently, minimize average client regret).² This can be cast as a special case of DirLat by defining the regret distances,

² Minimizing total regret is harder than MLP in the metric (V, c) . While an optimal MLP-solution for the metric (V, c) clearly yields an optimal solution to the minimum total-regret problem, this translation does not apply to near-optimal MLP-solutions. However, it is easy to see that an α -approximate solution to the minimum-total-regret problem is also an α -approximate MLP solution for the metric (V, c) .

$c_{u,v}^{\text{reg}} := c_{r,u} + c_{u,v} - c_{r,v}$, which form an asymmetric metric that we call the *regret metric* of (V, c) . We have not attempted to optimize the constant in Theorem 1, but it is rather large; we provide a substantially improved (and explicit) guarantee for the minimum-total-regret problem (Theorem 5) when (V, c) is a *symmetric* metric (i.e. $c_{u,v} = c_{v,u}$).

Our techniques. As noted above, our algorithm utilizes a natural time-indexed LP-relaxation for DirLat. Using standard scaling techniques, one may assume all $c_{u,v}$ distances are integers that are bounded by a polynomial in n (see Appendix A). Let $T = n \cdot \max_{u,v} c_{uv}$ and notice that T is bounded by a polynomial in n . Any Hamiltonian path in the metric $(V \cup \{r\}, c)$ has length at most T , so all nodes in the optimum solution are visited by time T .

For a directed graph $G = (N, E)$, and set $S \subseteq N$, let $\delta_G^{\text{in}}(S) := \{(u, v) \in E : u \in N - S, v \in S\}$ and $\delta_G^{\text{out}}(S) := \{(u, v) \in E : u \in S, v \in N - S\}$ denote respectively the edges entering and leaving S . Define $\delta_G(S) := \delta_G^{\text{in}}(S) \cup \delta_G^{\text{out}}(S)$. If the graph is clear from the context, we may omit the subscript G . We often identify an asymmetric metric $(V \cup \{r\}, c)$ with the complete directed graph over nodes $V \cup \{r\}$, with edge costs $c_{u,v}$ for distinct $u, v \in V \cup \{r\}$. For a path P and a node v on P , recall that $c_P(v)$ is the cost of the $r \rightsquigarrow v$ subpath of P . We begin with essentially the same time-indexed LP relaxation used in [22] for the undirected MLP, specifically (LP3) in their work. For $v \in V \cup \{r\}$ and $t \in [T]$, let $x_{v,t}$ be a variable indicating that we visit v at time *exactly* t , and let $z_{uv,t}$ indicate that we finish traversing edge (u, v) at time *exactly* t . Define $[T] := \{0, 1, \dots, T\}$.

$$\text{minimize :} \quad \sum_{v \in V, t \in [T]} t \cdot x_{v,t} \quad (\text{LP-Lat})$$

$$\text{subject to :} \quad \sum_{t \in [T]} x_{v,t} = 1 \quad \forall v \in V \quad (1)$$

$$\sum_{e \in \delta^{\text{in}}(S)} \sum_{t' \leq t} z_{e,t'} \geq \sum_{t' \leq t} x_{v,t'} \quad \forall v \in V, \{v\} \subseteq S \subseteq V, t \in [T] \quad (2)$$

$$x_{v,t} = \sum_{e \in \delta^{\text{in}}(v)} z_{e,t} \geq \sum_{e \in \delta^{\text{out}}(v)} z_{e,t+c_e} \quad \forall v \in V, t \in [T] \quad (3)$$

$$x, z \geq 0.$$

(We remark that the $z_{uv,t}$ variables above have a slightly different meaning from [22], wherein $z_{uv,t}$ indicated that t was traversed *by* time t . Also, we omit constraints (14) from [22], which encode that the length- t prefix of the optimal path has length at most t , as one can easily show they are implied by our slightly different approach.)

It is easy to check that an optimal solution P^* naturally corresponds to an integral solution to (LP-Lat) with the same cost as the latency of P^* . Constraints (2) admit an efficient separation oracle simply by checking for each $v \in V$ and $t \in T$ if the minimum $r - v$ cut has capacity at least $\sum_{t' \leq t} x_{v,t'}$ when using a capacity of $\sum_{t' \leq t} z_{e,t'}$ for each edge e .

Our proof of Theorem 1 proceeds by bucketing clients based on their fractional latencies, finding low-cost paths for these buckets, and stitching these paths together to form our final path. Our advantage over [12] comes from the fact that we guess the $O(\log T) = O(\log n)$ nodes v_i^* appearing at distances roughly 2^i along the optimum path P^* , plus their exact visiting times, ℓ_i^* , along P^* . We add constraints to (LP-Lat) to reflect these guesses. For each v_i^* , consider the nodes v that are at least, say, $2/3$ -visited before v_i^* but not $2/3$ -visited before v_{i-1}^* is visited: call this the *bucket* B_i for v_i^* . With a bit of modification, the restriction of (LP-Lat) to the times before ℓ_i^* is visited induces an LP solution with cost $O(2^i)$ for the natural ATSP LP relaxation that covers all $v \in B_i$ to an extent of at least $2/3$. That is, we get a solution to the following LP relaxation for ATSP for $\rho = 2/3$.

$$\begin{aligned}
& \text{minimize :} && \sum_{u,v} c_{u,v} \cdot x_{u,v} && (\text{LP-ATSP}_{\rho}) \\
& \text{subject to :} && x(\delta^{\text{out}}(v)) - x(\delta^{\text{in}}(v)) = \begin{cases} +1 & v = s \\ -1 & v = t \\ 0 & v \neq s, t \end{cases} && \forall v \in V \\
& && x(\delta(U)) \geq 2 \cdot \rho && \forall \emptyset \subsetneq U \subseteq V - \{s, t\} \\
& && x \geq 0.
\end{aligned}$$

The integrality gap when $\rho = 1$ was shown to be constant in [17]. We prove the following more-general result establishing a constant integrality gap for (LP-ATSP_{ρ}) for all $1/2 < \rho \leq 1$, which is one of our chief technical results. By an *LP-relative α -approximation algorithm* for (LP-ATSP_{ρ}) (or simply LP-relative approximation algorithm), we mean a polytime algorithm that returns an ATSP solution of cost at most $\alpha \cdot \text{OPT}_{\text{LP-ATSP}_{\rho}}$.

► **Theorem 2.** *There is an LP-relative $\frac{\psi}{2^{\rho-1}}$ -approximation algorithm for (LP-ATSP_{ρ}) , where ψ is some absolute constant (i.e., independent of the instance).*

We do not compute the exact value of ψ , or attempt to optimize it (favoring simplicity of presentation instead). It's precise value depends on the integrality gap for ATSP, which is known to be bounded by a constant [25, 28].

Using Theorem 2, we can obtain a path P_i for each bucket B_i , of cost $O(2^i)$ spanning the nodes of $\{r\} \cup B_i$. Our final path Q will be the concatenation of these P_i paths. To obtain Theorem 1, it suffices to show that the latency under Q of each node in B_i is $O(2^i)$. For the latter, while $c(P_i) = O(2^i)$, we also need a bound of $O(2^i)$ on the cost of stitching the last node of P_{i-1} to the first node after r on P_i . This is where guessing plays the most prominent role: we show that strengthening the LP with our guess ultimately implies the new edge used to stitch P_{i-1} to P_i also has cost $O(2^i)$, as required.

As an aside, complementing Theorem 2, we show that the dependence of the integrality gap on ρ stated in Theorem 2 is asymptotically correct, and this holds even if we strengthen (LP-ATSP_{ρ}) to require an in-flow of 1 for each $v \in V - \{s, t\}$ (but still have the relaxed cut constraints). This generalizes a similar result in [12] showing that the integrality gap of (LP-ATSP_{ρ}) is unbounded when $\rho = 1/2$.

► **Theorem 3.** *The integrality gap of (LP-ATSP_{ρ}) is at least $\frac{1}{2^{\rho-1}}$, for every $1/2 < \rho \leq 1$, and this holds even if we strengthen the LP with the constraints $x(\delta^{\text{in}}(v)) = 1$ for each $v \in V - \{s, t\}$.*

Our final result pertains to the minimum total-regret problem, for which we obtain a much-improved approximation guarantee (compared to Theorem 1). Recall that this is the special case of DirLat, where the metric is the regret metric of an *undirected* metric; in the sequel, we refer to this simply as a regret metric. Our improvement stems from the following improved and explicit integrality gap for (LP-ATSP_{ρ}) in regret metrics.

► **Theorem 4.** *There is an LP-relative $\alpha_{\rho}^{\text{reg}}$ -approximation algorithm for (LP-ATSP_{ρ}) in regret metrics, where $\alpha_{\rho}^{\text{reg}} := \frac{300}{42-12\sqrt{6}} \cdot \frac{1}{2^{\rho-1}} \approx \frac{23.8}{2^{\rho-1}}$.*

The proof of the above result is quite different from that of Theorem 2. It exploits the structure of regret metrics, and leverages and builds upon the insights and machinery developed in [13, 14] for this class of metrics. Theorem 4 leads to the following explicit approximation factor for DirLat in regret metrics.

► **Theorem 5.** *There is a quasi-polynomial time 397-approximation for DirLat in regret metrics.*

2.1 Related Work

Nagarajan and Ravi first studied DirLat and obtained an approximation guarantee of $n^{1/2+\epsilon}$ in time $n^{O(1/\epsilon)}$ for any constant $\epsilon > 0$ [20], which extends easily to an $O(\alpha' \cdot \log^{O(1)}(n))$ -approximation in quasi-polynomial time where (roughly speaking) α' is an upper bound on the integrality gap of the natural Held-Karp LP relaxation for ATSP. They also showed α' is bounded by $O(\sqrt{n})$. Friggstad, Salavatipour, and Svitkina improved the approximation guarantee for DirLat and the upper bound on the integrality gap for ATSP to $O(\log n)$ [12]. This is currently the best polynomial-time approximation for DirLat and no better quasi-polynomial time approximation was known before our work. If the metric is symmetric, constant-factor approximations are known. The first was given by Blum et al. [5], the best guarantee so far is a 3.59-approximation by Chaudhuri et al. [9].

Chakrabarty and Swamy [7], and Post and Swamy [22] studied LP relaxations for the undirected minimum latency problem. Using time-indexed LP relaxations, [22] obtain improved approximations for the multi-depot variant and also recover the 3.59-approximation for the single-vehicle version using an LP relaxation. Our work builds upon the ideas behind one of their LP relaxations.

The integrality-gap upper bound for ATSP has seen various improvements since [12], which have followed analogous improvements on the integrality gap, denoted α_{ATSP} , of the Held-Karp relaxation for ATSP, its more well-studied cousin. Friggstad et al. [11] show that the integrality gap is $O(\log n / \log \log n)$ by building upon ideas introduced in [3] who proved a similar bound for α_{ATSP} . Recently, [17] shows the integrality gap is in fact $O(1)$. Specifically, they show the gap is at most $4 \cdot \alpha_{\text{ATSP}} - 3$; combined with a breakthrough result of Svensson, Tarnawski, and Vegh [25], who showed $\alpha_{\text{ATSP}} = O(1)$, this yields an $O(1)$ upper bound on the integrality gap for ATSP. An even more recent development by Traub and Vygen shows that $\alpha_{\text{ATSP}} \leq 22$ [28], and Traub [27] has shown the integrality gap for ATSP is at most 43. The best lower bound known on α_{ATSP} is 2 [8].

The notion of regret has been proposed in the vehicle-routing literature (see, e.g., [24, 21]) as a more refined way of measuring client dissatisfaction than simply considering its waiting time. The underlying motivation is that since the shortest-path distance of a client from the depot is an inherent lower bound on its waiting time, it is more meaningful to measure the waiting time of a client *relative* to this lower bound. In symmetric metrics, two main regret-related problems have been investigated: finding a path (or a fixed number of paths) that minimizes maximum client regret; and finding the fewest number of bounded-regret paths to visit all clients. Constant-factor approximation algorithms are known for both problems (see [13] and the references therein). To our knowledge there is no prior work on finding provably near-optimal solutions for the total-regret (or equivalently average-regret) objective.

Outline of the paper. Section 3 presents the proofs of Theorems 1 and 5, assuming the LP-relative approximation algorithms provided by Theorems 2 and 4. Section 4 proves Theorem 2 and is concluded with the proof of Theorem 3. Finally, the proof of Theorem 4 is presented in Section 5.

3 An $O(1)$ -Approximation in Quasi-Polynomial Time

In this section, we assume Theorems 2 and 4 and use them to prove Theorems 1 and 5. By scaling (see Theorem 26 in Appendix A), we may assume distances are integers bounded by a polynomial in n and that $c_{u,v} \geq 1$ for distinct nodes u, v . We also let $T = n \cdot \max_{u,v \in V \cup \{r\}} c_{u,v}$,

which is an upper bound on the cost of any Hamiltonian path. We focus on a fixed optimal path P^* . Our algorithm starts by guessing the last node v_i^* visited by P^* at some time in the interval $[2^i, 2^{i+1})$ (if any) and its exact distance $\ell_i^* \in [T]$ for each $0 \leq i \leq \log_2 T = O(\log n)$. Let $v_i^* = \perp$ if no such node exists for this interval. For any i , we then know that no node is visited at any time in $[2^i, 2^{i+1})$ if $v_i^* = \perp$ and, if $v_i^* \neq \perp$, we also know no node is visited at a time in the interval $(\ell_i^*, 2^{i+1})$ so we mark these times as **forbidden**. Let $A = \{i : v_i^* \neq \perp\}$ be **admissible** buckets corresponding to intervals where the optimum visits at least one node. Let $1/2 < \rho \leq 1$ be a parameter we optimize later.

■ **Algorithm 1** Directed Latency: $O(1)$ -approximation in $n^{O(\log n)}$ time.

Input: asymmetric metric $(V \cup \{r\}, c)$ with integer distances at most T/n ; parameter $\rho \in (1/2, 1]$; an LP-relative α_ρ -approximation algorithm Alg for (LP-ATSP_ρ) .

Output: an r -rooted path P

-
- D1.** For every choice (guess) of $v_i^* \in V \cup \{\perp\}$ for each $0 \leq i \leq \log_2 T$ and $\ell_i^* \in [T]$ for each such i where $v_i^* \neq \perp$, perform the following steps. Let $F = \{t \in [T] : t \in [2^i, 2^{i+1}) \text{ where } v_i^* = \perp \text{ or } t \in (\ell_i^*, 2^{i+1}) \text{ where } v_i^* \neq \perp\}$ be the forbidden times for this guess (v^*, ℓ^*) and $A = \{i \in [0, \log_2 T] : v_i^* \neq \perp\}$ the admissible buckets.
- D1.1.** Obtain an optimal extreme point solution (x, z) to (LP-Lat) strengthened with the following additional constraints: 1) $x_{v_i^*, \ell_i^*} = 1$ for each $i \in A$ and 2) $x_{v,t} = 0$ for each $v \in V$ and $t \in F$. If the LP is infeasible, abort this guess of (v^*, ℓ^*) .
- D1.2.** For each $v \in V$, let $t(v) = t_\rho(v)$ be the minimum time such that $\sum_{t \leq t(v)} x_{v,t} \geq \rho$. For $i \in A$, let $B_i = \{v \in V : t(v) \in [2^i, 2^{i+1})\}$.
- D1.3.** For each $i \in A$, use algorithm Alg to obtain an $r - v_i^*$ path P_i spanning $\{r\} \cup B_i$.
- D1.4.** Let P^{v^*, ℓ^*} be the path obtained by concatenating the paths $\{P_i\}_{i \in A}$ in increasing order of i , and shortcutting past repeat occurrences of r .
- D2.** Return the best path P^{v^*, ℓ^*} found over all guesses where the strengthening of (LP-Lat) was feasible.
-

Let P^* be an optimum solution and consider the iteration where (v^*, ℓ^*) is consistent with P^* . Let (x, z) be an optimum LP solution for the strengthening of (LP-Lat) by the constraints in Step (11). Clearly this strengthened LP is feasible and the value of the solution (x, z) is at most OPT , the latency of P^* .

For each $v \in V$, note that $t(v)$ is well-defined by Constraints (1). Ultimately, we will show the path P^{v^*, ℓ^*} visits each $v \in V$ by time $O(t(v))$. We begin by showing this suffices to get a constant-factor approximation.

► **Lemma 6.** *Let P be a path and $c \geq 1$ be such that $c_P(v) \leq c \cdot t(v)$ for each $v \in V$. Then the latency of P is at most $\frac{c}{1-\rho} \cdot OPT$.*

Proof. Fix some $v \in V$. By definition of $t(v)$, we have $\sum_{t(v) \leq t \leq T} x_{v,t} \geq 1 - \rho$ which yields $t(v) \leq \frac{1}{1-\rho} \cdot \sum_{t(v) \leq t \leq T} t(v) \cdot x_{v,t} \leq \frac{1}{1-\rho} \cdot \sum_{t \in [T]} t \cdot x_{v,t}$. It follows that $\sum_v c_P(v) \leq \frac{c}{1-\rho} \cdot OPT$. ◀

3.1 Bounding the Latency of P^{v^*, ℓ^*}

In the remainder of the proof it is convenient to view a “time-expanded” graph G_T . The nodes are pairs (v, t) with $v \in V \cup \{r\}$ and $t \in [T]$ and an edge connects (u, t) to (v, t') if $c_{u,v} = t' - t$. Observe G_T is acyclic. We can then view $z_{e,t}$ as assigning values to edges of G_T : the edge $(u, t - c_{u,v}), (v, t)$ has value $z_{(u,v),t}$ and cost $c_{u,v}$.

We begin with some observations. The constraints of (LP-Lat) mean z constitutes one unit of $(r, 0)$ -preflow³ in G_T (i.e. a preflow with source vertex $(r, 0)$). Namely, Constraints (3) ensure preflow is satisfied at every vertex (v, t) of G_T apart from the “source” vertex $(r, 0)$. Let i' be the greatest index in A . Considering the LP constraints added in Step (11), we see $x_{v_{i'}, \ell_{i'}^*} = 1$ and $x_{v, t} = 0$ for all $t > \ell_{i'}^*$. Thus, z must be a flow with value 1 in G_T ending at $(v_{i'}, \ell_{i'}^*)$. Since the support of the flow z is acyclic in G_T and since one unit of flow passes through *every* (v_i^*, ℓ_i^*) node in G_T for each $i \in A$, no flow skips past node (v_i^*, ℓ_i^*) . That is, no edge $((u, t), (v, t'))$ in G_T supports any z -flow if $t < \ell_i^* < t'$ for some $i \in A$, nor does any edge $((u, t), (v, t'))$ support any z -flow if $t = \ell_i^*$ yet $u \neq v_i^*$ or $t' = \ell_i^*$ yet $v \neq v_i^*$ for some $i \in A$.

Next, we recall a famous splitting-off result by Mader. The following is a slight specialization of one such result.

► **Theorem 7** (Mader [18]). *Let $D = (V \cup \{s\}, A)$ be a directed, Eulerian multigraph such that the $u - v$ connectivity for every $u, v \in V$ is at least k . Then for every $(u, s) \in A$ there is some $(s, v) \in A$ such that in the graph $D' = (V \cup \{s\}, A - \{(u, s), (s, v)\} \cup \{(u, v)\})$, the $u - v$ connectivity for every $u, v \in V$ is still at least k .*

Using this, we show how to compute low-cost paths covering each bucket. Roughly speaking, we show that (LP-ATSP $_\rho$) restricted to $\{r\} \cup B_i$ with start node r and end node v_i^* has cost at most 2^{i+1} . Thus, step 13 would find a path starting at r and covering all B_i with cost at most $\alpha_\rho \cdot 2^{i+1}$ where we recall α_ρ denotes the approximation factor of the LP-relative approximation Alg.

► **Lemma 8.** *For each $i \in A$, we can compute a Hamiltonian $r - v_i^*$ path P_i in $G[\{r\} \cup B_i]$ with cost $\alpha_\rho \cdot 2^{i+1}$ in polynomial time.*

Proof. Let x' be a vector over edges of the metric given by $x'_{u,v} = \sum_{t < 2^{i+1}} z_{(u,v),t}$ for $u, v \in V \cup \{r\}$. From the observations above, the restriction of z to edges $((u, t), (v, t'))$ where $t < 2^{i+1}$ constitutes one unit of flow from $(r, 0)$ to (v_i^*, ℓ_i^*) in G_T , so x'_{uv} is then one unit of $r - v_i^*$ flow in the metric. Further, since the cost of an edge $((u, t - c_{u,v}), (v, t))$ is $c_{u,v}$ in G_T , the cost of this flow x' is exactly ℓ_i^* , which is at most 2^{i+1} .

Next we verify $x'(\delta(S)) \geq 2 \cdot \rho$ for each $S \subseteq V - \{v_i^*\}$ with $S \cap B_i \neq \emptyset$. Consider some $v \in S \cap B_i$. Constraint (2), the fact that $v \in B_i$, and the fact that $x_{v,t} = 0$ for $\ell_i^* < t < 2^{i+1}$ shows $x'(\delta^{\text{in}}(S)) = \sum_{e \in \delta(S)} \sum_{t < 2^{i+1}} z_{e,t} \geq \rho$. Since x' is an $r - v_i^*$ flow and $r, v_i^* \notin S$, then flow conservation shows $x'(\delta(S)) \geq 2 \cdot \rho$.

Much like in [2] for the PRIZE-COLLECTING TSP-PATH problem, one can use Theorem 7 to shortcut x' past nodes not in $B_i \cup \{r\}$ to get solution for (LP-ATSP $_\rho$) for in the graph $G[\{r\} \cup B_i]$ (with start node $s = r$ and end node $t = v_i^*$), also with cost at most 2^{i+1} . That is, we may assume x' is rational as z is a rational vector since it is part of an extreme point of an LP with rational coefficients. Let Δ be an integer such that the vector $\Delta \cdot x'$ is integral. Consider the graph G' with nodes $V \cup \{r\} \cup \{r'\}$ where r' is a new node. The edges of G' consist of $\Delta \cdot x'_{uv}$ copies of edge uv for each $u, v \in V \cup \{r\}$, and Δ edges from v_i^* to r' and also from r' to r (each having cost 0). Note the $r - u$ connectivity for each $u \in V$ is at least $\Delta \cdot \rho$. Note, the cost of all edges in G' is at most $\Delta \cdot 2^{i+1}$.

For each $v \in V - B_i$, we iteratively perform the splitting off procedure from Theorem 7 for $s = v$. The total cost of the edges does not increase by the triangle inequality (note the edges that are removed and added all lie in the metric over $V \cup \{r\}$), and the $r - u$

³ An s -preflow in a digraph (V, E) where $s \in V$ is an assignment $f : E \rightarrow \mathbb{R}_{\geq 0}$ such that $f(\delta^{\text{in}}(v)) \geq f(\delta^{\text{out}}(v))$ for each $v \in V - \{s\}$. The value of the preflow f is $f(\delta^{\text{out}}(s)) - f(\delta^{\text{in}}(s))$.

connectivity remains at least $\Delta \cdot \rho$ for each $u \in B_i$. After doing this for each $v \in V - B_i$, we are left with a multigraph of total edge cost no more than the total cost of all edges in G' . Further, if we remove all $v_i^* r'$ and $r' r$ edges, we still get the connectivity from r to any other $v \in B_i$ is at least $\Delta \cdot \rho$. If $k_{u,v}$ denotes the number of copies of uv in this new graph, setting $x''_{u,v} = k_{u,v}/\Delta$ for each $(u,v) \in G[\{r\} \cup B_i]$ yields a feasible LP solution for (LP-ATSP $_\rho$) in the metric graph over $B_i \cup \{r\}$ (with start node r and end node v_i^*) with cost at most 2^{i+1} .

From this, the optimal solution to (LP-ATSP $_\rho$) in $G[\{r\} \cup B_i]$ (starting at r and ending at v_i^*) has value at most 2^{i+1} . So Alg returns a Hamiltonian $r - v_i^*$ path P_i in $G[\{r\} \cup B_i]$ with cost at most $\alpha_\rho \cdot 2^{i+1}$. ◀

Next we bound the cost of stitching together the paths for the admissible buckets.

► **Lemma 9.** *Let P_i and $P_{i'}$ be two paths constructed in Step (13) for consecutive indices $i, i' \in A$. Let $u_{i'}$ be the first node on $P_{i'}$ after r and recall v_i^* is the last node of P_i . Then $c_{v_i^*, u_{i'}} \leq 2^{i'+1}$.*

Proof. Note that $u_{i'} \in B_{i'}$ means $t(u_{i'}) \in [2^{i'}, 2^{i'+1})$. Also, $x_{u_{i'}, t(u_{i'})} > 0$ by definition of $t(u_{i'})$. All units of z -flow in the acyclic graph G_T pass through (v_i^*, ℓ_i^*) and also through $(v_{i'}^*, \ell_{i'}^*)$. So the restriction of z to edges $((u, t), (v, t'))$ in G_T with $\ell_i^* \leq t \leq t' \leq \ell_{i'}^*$ constitutes one unit of $(v_i^*, \ell_i^*) - (v_{i'}^*, \ell_{i'}^*)$ flow that supports $(u_{i'}, t(u_{i'}))$. Therefore, a path decomposition of this restriction of z includes $(u_{i'}, t(u_{i'}))$ on some path. Any such path has cost exactly $\ell_{i'}^* - \ell_i^* \leq 2^{i'+1}$. By the triangle inequality, $c_{v_i^*, u_{i'}} + c_{u_{i'}, v_{i'}^*} \leq 2^{i'+1}$. ◀

Next, we bound the latency of each $v \in V$ along the final path P^{v^*, ℓ^*} obtained by concatenating the P_i paths for increasing indices $i \in A$ and shortcutting past all but the first occurrence of r .

► **Lemma 10.** $c_{P^{v^*, \ell^*}}(v) \leq 4(\alpha_\rho + 1) \cdot t(v)$ for any $v \in V$.

Proof. Consider any $v \in V$ and say it lies on P_i . To reach v along P^{v^*, ℓ^*} , we traverse paths $P_{i'}$ for $i' < i$ plus the “stitching” edges $v_{i'}^* u_{i'}^*$ for consecutive indices $i', i'' \in A$, $i'' \leq i$. By Lemma 8 and Lemma 9, the latency of v along P^{v^*, ℓ^*} can be bounded by $\sum_{i' \in A, i' \leq i} \alpha_\rho \cdot 2^{i'+1} + \sum_{i' \in A, i' \leq i} 2^{i'+1} \leq (\alpha_\rho + 1) \cdot \sum_{i'=0}^i 2^{i'+1} \leq 4(\alpha_\rho + 1) \cdot 2^i \leq 4(\alpha_\rho + 1) \cdot t(v)$. ◀

Proof of Theorem 1. We set $\rho = 2/3$, and note that Theorem 2 yields an LP-relative $\alpha_{2/3}$ -approximation algorithm, where $\alpha_{2/3} = O(1)$. The proof of Theorem 1 then follows readily from Lemmas 6 and 10 and the fact that T is bounded by a polynomial in n . ◀

We remark that even with the improved bound of $\alpha \leq 22$ from [28], our approach yields an approximation ratio in the thousands. As noted earlier, we obtain a much-better guarantee for the special case of regret metrics, i.e., the minimum-total-regret problem.

Proof of Theorem 5. First, we note that a worse approximation ratio follows by choosing $\rho = 0.75$: this yields an approximation ratio α_ρ of at most 47.6 for the LP-relative algorithm in Theorem 4 for regret metrics, which combined with Lemmas 10 and 6 (and choosing ϵ sufficiently small in Theorem 26) yields a 778-approximation.

The better guarantee stated in the theorem follows by choosing the best ρ tailored for the given instance. (Note that there are only polynomially many combinatorially-distinct choices of ρ , and we can simply try all of these to pick the best ρ .) We analyze this by choosing a random ρ and bounding the expected latency incurred; this is similar to the use of random α -points in scheduling algorithms (see, e.g., [23]).

For $v \in V$, recall that $t_\rho(v)$ is the minimum time for $\sum_{t \leq t(v)} x_{v,t} \geq \rho$. Define $\text{LP}_v := \sum_{t \in [T]} t x_{v,t}$. The key is to realize that $\int_0^1 t_\rho(v) d\rho = \text{LP}_v$, and leverage this in place of the coarse bound $t_\rho(v) \leq \frac{\text{LP}_v}{1-\rho}$ used earlier (in Lemma 6). The approximation factor α_ρ given by Theorem 4 is of the form $\frac{c}{2\rho-1}$, where $c = 23.8$. We choose a random ρ from $(1/2, 1]$ according to the density function $8(x - 1/2)$. The expected latency incurred by a node v is then at most

$$\int_{1/2}^1 4 \left(\frac{c}{2\rho-1} + 1 \right) t_\rho(v) \cdot 8(\rho - 1/2) d\rho \leq 16c \cdot \int_{1/2}^1 t_\rho(v) d\rho + 16 \int_{1/2}^1 t_\rho(v) d\rho \leq 16(c+1) \cdot \text{LP}_v.$$

Thus, the expected total latency is at most $16(23.8 + 1) \cdot \text{OPT} \leq 397 \cdot \text{OPT}$. ◀

4 Bounding the Integrality Gap of (LP-ATSPP $_\rho$)

Consider nodes V with two distinguished $s, t \in V$ ($s \neq t$) and asymmetric metric distances $c_{u,v}$ between points of V . We consider (LP-ATSPP $_\rho$) for the Asymmetric TSP Path problem where the goal is to find the cheapest Hamiltonian $s - t$ path. As mentioned earlier, the integrality gap is unbounded if $\rho \leq 1/2$ [12], so we focus on the case $1/2 < \rho \leq 1$. As in [17], we start with the dual of (LP-ATSPP $_\rho$).

$$\begin{aligned} \text{maximize :} \quad & z_t - z_s + \sum_U 2\rho \cdot y_U && (\text{DUAL}_\rho) \\ \text{subject to :} \quad & z_v - z_u + \sum_{U:uv \in \delta(U)} y_U \leq c_{u,v} && \forall u, v \\ & y \geq 0. \end{aligned}$$

Naturally, our proof borrows many steps from Köhne, Traub, and Vygen [17] but there are additional challenges we have to work through in this more general setting.

For a vector x over the edges E of the directed metric (when viewed as a complete, directed graph), let $\text{supp}(x) = \{uv \in E : x_{u,v} > 0\}$. Similarly, for a vector y over cuts of the metric let $\text{supp}(y) = \{\emptyset \subsetneq S \subseteq V - \{s, t\} : y_S > 0\}$. From now on, we focus on the graph $G = (V, \text{supp}(x))$. The proofs of Propositions 11, 12, and 14 are very similar to proofs in [17] and are omitted or just sketched in this paper.

► **Proposition 11.** *Given any optimal dual solution (y, z) , one can find an optimal dual solution (y', z) with $\text{supp}(y')$ being laminar in polynomial time.*

In other words, we can modify y to be laminar without changing z using efficient uncrossing techniques. The proof is exactly the same as the proof in [17] essentially because the set of feasible solutions to (DUAL $_\rho$) does not change if we select different values for ρ .

The next proposition is almost identical to one in [17], but we omit the case $U = V$ in the statement. In fact, the result may not be true for this case $U = V$, we handle that separately below.

► **Proposition 12.** *Let x be an optimum primal solution and let $G = (V, \text{supp}(x))$. For any $U \subseteq V - \{s, t\}$ with $x(\delta(U)) = 2\rho$, any topological ordering U_1, \dots, U_ℓ of the strongly connected components of $G[U]$ satisfies:*

- $\delta^{\text{in}}(U_1) = \delta^{\text{in}}(U)$,
- $\delta^{\text{out}}(U_\ell) = \delta^{\text{out}}(U)$, and
- $x(\delta^{\text{out}}(U_i)) = x(\delta^{\text{in}}(U_{i+1}))$ for any $1 \leq i < \ell$.

We sketch the proof of Proposition 12 so the reader is assured it holds, though the proof is essentially the same.

Proof sketch. Because U is a tight set, $x(\delta^{\text{in}}(U)) = \rho$. Further, $x(\delta^{\text{in}}(U_1)) \geq \rho$. All edges in $\text{supp}(x)$ entering $\delta(U_1)$ must lie in $\delta^{\text{in}}(U)$ because U_1 is the first node in the topological ordering. Thus, $\rho = x(\delta^{\text{in}}(U)) \geq x(\delta^{\text{in}}(U_1)) \geq \rho$, so equality must hold throughout and $\delta^{\text{in}}(U) = \delta^{\text{in}}(U_1)$ as we are working in the support of x . A similar statement shows $\delta^{\text{out}}(U_\ell) = \delta^{\text{out}}(U)$.

For $i > 1$ we note $\delta^{\text{in}}(U_i) \subseteq \delta^{\text{in}}(U) \cup \bigcup_{j < i} \delta^{\text{out}}(U_j)$ simply because the U_j are topologically ordered. Inductively, we have $x(\delta^{\text{out}}(U_{i-1})) = \rho$ and each edge in $\delta^{\text{in}}(U) \cup \bigcup_{j < i-1} \delta^{\text{out}}(U_j)$ is already proven to lie in $\delta^{\text{in}}(U_{j'})$ for some $j' < i$. So we see $\delta^{\text{in}}(U_i) \subseteq \delta^{\text{out}}(U_{i-1})$ and, thus,

$$\rho = x(\delta^{\text{in}}(U_{i-1})) = x(\delta^{\text{out}}(U_{i-1})) \geq x(\delta^{\text{in}}(U_i)) \geq \rho.$$

So, again, equality must hold throughout. \blacktriangleleft

We use a different observation to address the case $U = V$ that was omitted from Proposition 12. Intuitively, we show that it is still possible to buy a cheap set of edges to chain the strongly-connected components of G in sequence but the cost of these edges does increase relative to OPT_{LP} as $\rho \rightarrow 1/2$.

► Proposition 13. *In any topological ordering U_1, \dots, U_ℓ of the strongly connected components of G , for each $1 \leq i < \ell$ there is some edge $(u, v) \in \delta^{\text{out}}(U_i) \cap \delta^{\text{in}}(U_{i+1})$ with $c_{u,v} \leq \frac{1}{2\rho-1} \cdot \sum_{u',v' \in \delta^{\text{out}}(U_i) \cap \delta^{\text{in}}(U_{i+1})} c_{u',v'} x_{u',v'}$.*

Proof. This is easy for $i = 1$ and $i = \ell - 1$. For example, we have $x(\delta^{\text{in}}(U_2)) \geq \rho$ and all edges from $\delta^{\text{in}}(U_2)$ lie in $\delta^{\text{out}}(U_1)$. Thus, $x(\delta^{\text{out}}(U_1) \cap \delta^{\text{in}}(U_2)) \geq \rho$ so the cheapest edge in $\delta^{\text{out}}(U_1) \cap \delta^{\text{in}}(U_2)$ has cost at most $\frac{1}{\rho} \cdot \sum_{uv \in \delta^{\text{out}}(U_1) \cap \delta^{\text{in}}(U_2)} c_{u,v} x_{u,v}$. We finish by observing $1/\rho \leq 1/(2\rho - 1)$ as $\rho \leq 1$. A similar argument works for $i = \ell - 1$, so we now assume $1 < i < \ell - 1$.

We quickly introduce notation. For an index $1 \leq j \leq \ell$ let $U_{\leq j} = \cup_{1 \leq j' \leq j} U_{j'}$ and $U_{\geq j} = \cup_{j \leq j' \leq \ell} U_{j'}$. Let $\delta(X; Y)$ denote $\{uv \in \text{supp}(x) : u \in X, v \in Y\}$ for $X, Y \subseteq V$. With this notation, let $a = x(\delta(U_i; U_{i+1}))$, $b = x(\delta(U_i; U_{\geq i+2}))$, $c = x(\delta(U_{\leq i-1}; U_{i+1}))$, and $d = x(\delta(U_{\leq i-1}; U_{\geq i+1}))$. We have $a + b + c + d = x(\delta^{\text{out}}(U_{\leq i})) = 1$ as $\delta^{\text{out}}(U_{\leq i})$ is the disjoint union of the sets defining a, b, c, d . On the other hand, $\rho \leq x(\delta^{\text{out}}(U_i)) = a + b$ and $\rho \leq x(\delta^{\text{in}}(U_{i+1})) = a + c$. Therefore, $2\rho - 1 \leq (a + b) + (a + c) - (a + b + c + d) \leq a$ so $x(\delta^{\text{out}}(U_i) \cap \delta^{\text{in}}(U_{i+1})) \geq 2\rho - 1$. So the cheapest edge $(u, v) \in \delta^{\text{out}}(U_i) \cap \delta^{\text{in}}(U_{i+1})$ has $c_{u,v} \leq \frac{1}{2\rho-1} \cdot \sum_{(u',v') \in \delta^{\text{out}}(U_i) \cap \delta^{\text{in}}(U_{i+1})} c_{u',v'} x_{u',v'}$. \blacktriangleleft

► Proposition 14. *Let G be the support graph of an optimum solution x to (LP-ATSP_ρ) and (y, z) an optimum dual with $\text{supp}(y)$ laminar. For any $U \in \text{supp}(y) \cup \{V\}$ and any $u, w \in U$ with w being reachable from u in $G[U]$, there is a $v - w$ path in $G[U]$ that crosses each set $U' \in \text{supp}(y)$ at most twice for $U' \subsetneq U$.*

Again, the proof is the same as that in [17] which only relies on Proposition 12 for $U \in \text{supp}(y)$ (i.e. not on the case $U = V$ that we omitted from the proposition in our setting). We sketch the argument briefly to ensure the reader this still holds with the omission of $U = V$ from Proposition 12.

Proof. Consider any $u - w$ path P contained in $G[U]$. Suppose $U' \in \text{supp}(y)$ is maximal among all such sets where P re-enters U' after it exits U' . Let a be the first node of P in U' and b the last node of P in U' (it could be $a = u$ or $b = v$). Inductively, replace the $a - b$ portion of P with an $a - b$ path in $G[U']$ that enters and leaves every set $U'' \in \text{supp}(y)$ at most once for $U'' \subsetneq U'$. Repeat for all such maximal $U' \in \text{supp}(y)$. \blacktriangleleft

4.1 Constructing the Path

Let OPT_{LP} denote the optimum solution value to (LP-ATSP $_{\rho}$). Recall we let α denote an upper bound on the integrality gap of the standard Held-Karp relaxation for ATSP. We will prove the following lemma later.

► **Lemma 15.** *An optimal dual solution (y, z) with $\text{supp}(y)$ being laminar and $z_s - z_t \leq \frac{1}{2\rho-1} \cdot OPT_{LP}$ can be computed in polynomial time.*

Using this, we now turn to the main result of this section. Note, we are choosing simplicity in presentation over optimizing the constants in the guarantee.

Proof of Theorem 2. Complementary slackness ensures $x(\delta(U)) = 2\rho$ for each $U \in \text{supp}(y)$. Consider the edge support graph $G = (V, \text{supp}(x))$. Modify G to get an ATSP instance H by adding a new node \bar{v} and edges (t, \bar{v}) with cost OPT_{LP} and (\bar{v}, s) with cost 0.

It is easy to check that setting

$$x'_{u,v} = \begin{cases} \frac{1}{\rho} & \text{if } (u, v) \in \{(t, \bar{v}), (\bar{v}, s)\} \\ \frac{x_{u,v}}{\rho} & \text{otherwise} \end{cases}$$


yields a feasible solution for the ATSP-Circuit relaxation from [25] in instance H with cost $\frac{2}{\rho} OPT_{LP}$. Using [25], we can find a circuit W spanning all nodes in H with cost at most $\frac{2\alpha}{\rho} OPT_{LP}$ in polynomial time. This circuit must use the (t, \bar{v}) edge at least once as it visits \bar{v} . By deleting occurrences of (t, \bar{v}) and (\bar{v}, s) , we get $s - t$ walks W_1, \dots, W_k in G that collectively span all nodes in V with $\sum_j c(W_j) \leq \frac{2\alpha}{\rho} \cdot OPT_{LP} \leq 4\alpha \cdot OPT_{LP}$. We also point out $k \leq 4\alpha$ because in removing the k edges incident to \bar{v} to get the walks W_i , we removed a total edge cost of $k \cdot OPT_{LP}$ from a circuit whose cost is at most $4\alpha \cdot OPT_{LP}$, so $k \leq 4\alpha$. The walks W_1, \dots, W_ℓ are depicted in the top of Figure 1.

Let U_1, \dots, U_ℓ be the strongly connected components of the support graph G . For each U_i , let $\mathcal{W}_i = \{j : W_j \text{ visits a node in } U_i\}$ and note $|\mathcal{W}_i| \leq k$. Unlike the case $\rho = 1$ in [17], it could be that $j \notin \mathcal{W}_i$ for some U_i and W_j . For each $1 \leq i \leq \ell$ and each $j \in \mathcal{W}_i$, let $R_{i,j}$ denote the restriction of W_j to U_i . Now, if some W_j enters U_i , then once it leaves it cannot re-enter because U_i is a strongly connected component of G . So $R_{i,j}$ is a single walk for each $j \in \mathcal{W}_i$. For such (i, j) , let u_j^i and v_j^i be the first and last nodes of W_j in U_i .

Order \mathcal{W}_i as $j_1 < j_2 < \dots < j_{|\mathcal{W}_i|}$. By Proposition 14 and the fact each U_i is a strongly connected component, we can find paths P_{i,j_m} for $j_m \in \mathcal{W}_i$ from $v_{j_m}^i$ to $u_{j_{m+1}}^i$ (or u_1^i if $m = |\mathcal{W}_i|$) where $P_{i,j}$ enters and exits each $U' \in \text{supp}(y)$ with $U' \subsetneq U_i$ at most once and does not cross any other set in $\text{supp}(y)$. Then, for each i we get a circuit C_i spanning all nodes of U_i by adding the paths $P_{i,j}$ for $j \in \mathcal{W}_i$ to the walks $R_{i,j}$.

By Proposition 13, for each $1 \leq i < \ell$ there are edges $u'_i v'_{i+1} \in \delta^{\text{out}}(U_i) \cap \delta^{\text{in}}(U_{i+1})$ with cost at most $\frac{1}{2\rho-1}$ times the fractional cost of edges in $\delta^{\text{out}}(U_i) \cap \delta^{\text{in}}(U_{i+1})$. Also, say $v'_1 = s$ and $u'_\ell = t$. By fully traversing each C_i starting at v'_i and then continuing to follow it again to reach u'_i , we get $v'_i - u'_i$ walks W'_i spanning U_i . The final path P we output is the concatenation of the walks $W'_1, W'_2, \dots, W'_\ell$. Let $S = \{v'_i u'_{i+1} : 1 \leq i < \ell\}$ be the edges used to “stitch” these walks W'_i together. The bottom of Figure 1 depicts the final construction of P .

To bound the cost of P , first observe $c(S) \leq \frac{1}{2\rho-1} OPT_{LP}$ as the sets $\delta^{\text{out}}(U_i) \cap \delta^{\text{in}}(U_{i+1})$ are disjoint for $1 \leq i < \ell$. To bound the cost of the cycles C_i , we define a modified cost $c^y_{uv} = \sum_{U:uv \in \delta(U)}$ and observe $c(Q) = z_v - z_u + c^y(Q)$ for any $u - v$ path Q (the z -values for internal nodes of Q cancel).



■ **Figure 1**

Top: A depiction of $s - t$ walks W_1, \dots, W_k with $k = 4$ and the strongly-connected components U_1, \dots, U_ℓ with $\ell = 5$.

Bottom: The solid edges are the restrictions of the walks W_i to the strongly-connected components: these are the $R_{i,j}$ walks. The thin dashed edges in each component are the paths $P_{i,j}$ that stitch these $R_{i,j}$ to form a circuit C_i over the strongly connected component U_i . Finally, the dashed edges between components are the edges in S obtained from Proposition 13. The final path P is obtained by visiting the U_i consecutively using these dashed edges, where each visit traverses C_i fully and then travels to the start of the edge exiting U_i .

By complementary slackness, $c_{u,v} = z_v - z_u + c_{uv}^y$ for each $uv \in \text{supp}(x)$. Each C_i was formed by stitching together endpoints of $R_{i,j}$ using paths $P_{i,j}$. Each $P_{i,j}$ crosses each $U' \in \text{supp}(y), U' \subsetneq U_i$ at most twice and does not cross any set in $\text{supp}(y)$ not contained in U_i . Further, no two $P_{i,j}, P_{i',j'}$ paths for $i \neq i'$ can cross the same $U' \in \text{supp}(y)$ because the two paths are contained in different components of G .

Each $U' \in \text{supp}(y)$ is crossed by at most k paths of the form $P_{i,j}$ meaning $\sum_{i,j} c^y(P_{i,j}) \leq \sum_{i,j} z_{v_j^i} - z_{u_j^i} + 2k \cdot \sum_U y_U$. We also have $c^y(R_{i,j}) = z_{u_j^i} - z_{v_j^i} + c(R_{i,j})$. Therefore, $\sum_i c^y(C_i) = \sum_i \sum_{j \in \mathcal{W}_{i,j}} c^y(P_{i,j}) + c^y(R_{i,j}) \leq 2k \sum_U y_U + \sum_{i,j \in \mathcal{W}_i} c(R_{i,j}) \leq 2k \sum_U y_U + \sum_j c(W_j)$ (the z terms for the endpoints of the $R_{i,j}$ cancel out in the first inequality).

But $c(C) = c^y(C)$ for any cycle C because, again, the z -terms cancel out. So

$$\begin{aligned}
 c(P) &\leq c(S) + 2 \cdot \sum_i c(C_i) &\leq \frac{OPT_{LP}}{2^{\rho-1}} + 2 \sum_{i=1}^k c(W_i) + 2k \sum_U y_U \\
 &\leq \frac{OPT_{LP}}{2^{\rho-1}} + 4\alpha \cdot OPT_{LP} + 2k \sum_U y_U &\leq O(1) \cdot \frac{1}{2^{\rho-1}} \cdot OPT_{LP} + \frac{k}{\rho} (OPT_{LP} + z_s - z_t) \\
 &\leq \frac{O(1)}{2^{\rho-1}} \cdot OPT_{LP} + \frac{k}{\rho} \cdot (z_s - z_t).
 \end{aligned}$$

Here, $O(1)$ refers to some constant that is independent of ρ where we also recall k is bounded by a constant as well. Using Lemma 15 to bound $z_s - z_t$ finishes the proof. ◀

4.2 Bounding $z_s - z_t$: Proof of Lemma 15

We prove Lemma 15 to finish the proof of Theorem 2. Our approach is more direct than [17], they used an argument that shifts LP weight around to show that $y_U > 0$ implies U is not an $s - t$ separator in the support graph $G = (V, \text{supp}(x))$. We establish this fact using complementary slackness applied to the LP used to find the optimal solution to DUAL_ρ with minimum possible $z_s - z_t$. We comment that their proof could also be adapted to show what we want, we are presenting this alternative proof because we feel it is more naturally motivated: we already want to minimize $z_s - z_t$ among all optimal duals so it is natural to ask what complementary slackness gives for $y_U > 0$.

Let x be an optimal primal solution to LP-ATSP_ρ . Note that if we restricted the variables of (LP-ATSP_ρ) and the constraints of (DUAL_ρ) to $\text{supp}(x)$ then x and (y, z) remains optimal. For any feasible solution (y, z) to (DUAL_ρ) , we know $z_t - z_s \leq \text{OPT}_{LP}$ because $y \geq 0$. So the following LP is bounded. Note, we first solved (LP-ATSP_ρ) to compute OPT_{LP} which is then a fixed value (not a variable) in $\text{DUAL}_{\rho-Z}$ below.

$$\text{maximize :} \quad z_t - z_s \quad (\text{DUAL}_{\rho-Z})$$

$$\text{subject to :} \quad z_t - z_s + \sum_{\emptyset \subsetneq U \subseteq V - \{s, t\}} 2\rho \cdot y_U \geq \text{OPT}_{LP} \quad (4)$$

$$z_v - z_u + \sum_{U: uv \in \delta(U)} y_U \leq c_{u,v} \quad \forall u, v \in \text{supp}(x) \quad (5)$$

$$y \geq 0.$$

The second constraint asserts (y, z) is a feasible solution for (DUAL_ρ) , so the first constraint then asserts it is an optimal solution for DUAL_ρ . In fact, in any feasible solution the first constraint must hold with equality. We prove $z_s - z_t \leq \frac{1}{2\rho-1} \cdot \text{OPT}_{LP}$ for an optimal solution (y, z) to $(\text{DUAL}_{\rho-Z})$. With this, we finish the proof of Lemma 15 by simply noting that Proposition 11 shows we can uncross the support of y while leaving z unchanged.

The LP that is dual to $(\text{DUAL}_{\rho-Z})$ has a variable κ for Constraint (4) of $(\text{DUAL}_{\rho-Z})$ and new variables x'_{uv} for each instance uv of Constraint (5).

$$\text{minimize :} \quad \sum_{uv \in \text{supp}(x)} c_{u,v} \cdot x'_{uv} - \text{OPT}_{LP} \cdot \kappa$$

$$\text{subject to :} \quad x'(\delta^{\text{out}}(v)) - x'(\delta^{\text{in}}(v)) = \begin{cases} 1 + \kappa & v = s \\ -1 - \kappa & v = t \\ 0 & v \neq s, t \end{cases} \quad \forall v \in V$$

$$x'(\delta(U)) \geq 2\rho \cdot \kappa \quad \forall \emptyset \subsetneq U \subseteq V - \{s, t\}$$

$$x', \kappa \geq 0.$$

► **Lemma 16.** *In an optimal solution (y, z) to $\text{DUAL}_{\rho-Z}$, if $y_U > 0$ then there is an $s - t$ path in the graph $G[V - U]$.*

Proof. Let x' be an optimal solution to the dual of $(\text{DUAL}_{\rho-Z})$. Then $y_U > 0$ implies $x'(\delta(U)) = 2\rho \cdot \kappa$ so, by flow conservation, $x'(\delta^{\text{in}}(U)) = \rho \cdot \kappa$.

On the other hand, x' constitutes an $s - t$ flow of value $1 + \kappa$. Consider a decomposition of x' into paths and cycles. The total weight of paths that do not enter U is at least $1 + \kappa - \rho \cdot \kappa = 1 + (1 - \rho) \cdot \kappa > 0$. Thus, there is an $s - t$ path in G that does not pass through U . ◀


Continuing as in [17], let U_1, \dots, U_k be the maximal sets in $\text{supp}(y)$. In the graph G' obtained by contracting each U_i , we have by Lemma 16 that for each contracted node U_i there is an $s - t$ path in G' that avoids U_i . By a variant of Menger's Theorem (Lemma 9 in [17]), there are node-disjoint $s - t$ paths P_1, P_2 in G' . Consider the edges of P_1 and P_2 in G . For any U_i , at most one of P_1 or P_2 enters (and exits) U_i . Suppose it is the case that one of them $\bar{P} \in \{P_1, P_2\}$ enters U_i . Let u, v be the first and last nodes of \bar{P} as it passes through U_i . By Proposition 14, we can find a $u - v$ path in $G[U_i]$ that crosses each $U' \in \text{supp}(y)$ contained in U at most twice, and does not cross any other set in $\text{supp}(y)$. Add these edges to \bar{P} .

Do this for each U_i that is entered by some $\bar{P} \in \{P_1, P_2\}$. We get paths P'_1, P'_2 using only edges in $\text{supp}(x)$ that, collectively, cross each set in $\text{supp}(y)$ at most twice. Thus, $0 \leq c(P_1) + c(P_2) = c^y(P_1) + c^y(P_2) + 2 \cdot (z_t - z_s) \leq 2 \cdot \sum_{U \in \text{supp}(y)} y_U + 2 \cdot (z_t - z_s)$. Multiplying the terms in this bound by ρ and then subtracting $(2\rho - 1) \cdot (z_t - z_s)$ from both sides, we see $(2\rho - 1) \cdot (z_s - z_t) \leq \sum_{U \in \text{supp}(y)} 2\rho \cdot y_U + z_t - z_s = \text{OPT}_{LP}$. ◀

4.3 A Bad Example for (LP-ATSP $_\rho$)

We show that the dependence on the factor $\frac{1}{2\rho-1}$ in our analysis of the integrality gap of (LP-ATSP $_\rho$) is asymptotically tight.

Proof of Theorem 3. Consider the following metric depicted in Figure (2), which is inspired from the example showing the integrality gap is unbounded if $\rho = 1/2$ from [12]. The solid edges have cost 0 and the dashed edges have cost 1. The cost of all other edges not depicted is the shortest path distance in this graph (using a cost of 1 if there is no path in this graph). The number beside each edge uv indicates the value of $x_{u,v}$. It can be easily check that this is a feasible solution for (LP-ATSP $_\rho$) even if we added the constraints $x(\delta^{\text{in}}(v)) = 1$ for each $v \in V - \{s, t\}$. An optimal integral solution must use an edge with cost 1, yet this LP solution only has cost $2\rho - 1$ so the integrality gap of (LP-ATSP $_\rho$) is at least $\frac{1}{2\rho-1}$. ◀



■ **Figure 2** The bad integrality gap example for LP-ATSP $_\rho$.

5 An Improved Integrality Gap Bound for (LP-ATSP $_\rho$) in Regret Metrics

Let V be nodes and $s, t \in V$ be the start and end points. Let c be *symmetric* metric distances $c_{u,v} \geq 0$. For each $u, v \in V$, let $c_{u,v}^{\text{reg}} = c_{r,u} + c_{u,v} - c_{r,v}$ be the *regret* metric induced by c . It is convenient to consider a complete directed graph over V where for distinct $u, v \in V$ we have $c_{u,v} = c_{v,u}$ yet (u, v) and (v, u) are themselves distinct edges: the *bidirected* variant of the natural undirected graph associated with (V, c) . The following observations about regret metrics can be found in [13].

► **Observation 17.** *If c is a metric (asymmetric or symmetric) then c^{reg} is an asymmetric metric. For any $u, v \in V$ and any $u - v$ path P , $c(P) = c^{\text{reg}}(P) + c_{u,v}$. For any cycle C , $c(C) = c^{\text{reg}}(C)$.*

We consider integrality gap bounds for (LP-ATSP $_{\rho}$) when the metric is a regret metric. In [14], it was shown the integrality gap bound is 2 in the standard case $\rho = 1$ and that this is tight. For the purpose of getting better approximations for DirLat in regret metrics (i.e. the problem of minimizing the average time a node v waits *in excess* of their shortest path distance $c_{r,v}$ from the depot), we give explicit integrality gap bounds for the more general case $1/2 < \rho \leq 1$.

Note, in the case $\rho = 1$ that the analysis from [14] produces a stronger result. But the analysis does not extend in any clear way to the case $\rho < 1$. We begin by recalling the following structural result by Bang-Jensen et al about decomposing preflows into branchings [4], which was made efficient by Gabow [15] (see also [22]).

► **Theorem 18** (Bang Jensen et al. [4], Gabow [15], Post and Swamy [22]). *Let $D = (\{r\} \cup V, A)$ be a directed graph and $x \in \mathbb{Q}_{\geq 0}^A$ be a preflow. Let $\lambda_v := \min_{\{v\} \subseteq S \subseteq V} x(\delta^{\text{in}}(S))$ be the $r - v$ connectivity in D under capacities $\{x_a\}_{a \in A}$. Let $K > 0$ be rational. We can obtain out-branchings B_1, \dots, B_q rooted at r , and rational weights $\gamma_1, \dots, \gamma_q \geq 0$ such that $\sum_{i=1}^q \gamma_i = K$, $\sum_{i:q \in B_i} \gamma_i \leq x_a$ for all $a \in A$, and $\sum_{i:v \in B_i} \gamma_i \geq \min\{K, \lambda_v\}$ for all $v \in V$. Moreover, such a decomposition can be computed in time that is polynomial in $|V|$ and the bit complexity of K and x .*

We require a definition and results from [13], some of which are adaptations of concepts in [6].

► **Definition 19.** *Let P be a path starting at s . For each $uv \in P$, say uv is **red** on P if there are nodes x, y on the $s - u$ portion of P_i and $v - t$ portion of i , respectively, such that $c_{r,x} \geq c_{r,y}$. For each $v \in P$, let $\text{red}(v, P)$ be the maximal subset of red edges of the subpath of P containing v . Note, $\text{red}(v, P)$ could be empty if v is not incident to a red edge. The **red intervals** of P are the maximal subpaths of its red edges.*

Intuitively, the red edges are part of intervals of P that do not make progress toward reaching t . Their total c^{reg} -costs can be shown to be comparable to their total c -costs, which is formalized as follows.

► **Lemma 20** (Blum et al [6]). *For any $s - t$ path P , $\sum_{uv \text{ red on } P} c_{u,v} \leq \frac{3}{2} c^{\text{reg}}(P)$.*

Further, if we were to keep at most one node from each maximal red interval of edges and shortcut past the other nodes, the resulting path $s = v_0, v_1, \dots, v_k = t$ has $c_{r,v_i} < c_{r,v_{i+1}}$. So the union of any collection of paths that are shortcut in such a way forms an acyclic graph.

Now, a solution to (LP-ATSP $_{\rho}$) can be viewed as a preflow of value 1 rooted at s with $\lambda_v \geq \rho$ for each $v \in V - t$ and $\lambda_t = 1$. From this observation, we round a solution using techniques from [13]. The full description is in Algorithm 2. Here, $1/2 < \delta < \rho$ is some parameter we set later to optimize the performance of the algorithm.

► **Lemma 21.** *The paths P_i from Step 2 satisfy $\sum_i \gamma_i \cdot c^{\text{reg}}(P_i) \leq 2 \cdot \text{OPT}_{LP}$.*

Proof. In [13], it is observed for any $s - t$ path P that $c^{\text{reg}}(P) = c(P) - c_{s,t}$ and that $c(C) = c^{\text{reg}}(C)$ for any cycle C . Thus, as x is an $s - t$ flow with value 1 we have $\text{OPT}_{LP} = \sum_{uv} c_{u,v}^{\text{reg}} x_{u,v} = (\sum_{uv} c_{u,v} x_{u,v}) - c_{s,t}$. This can be seen by, say, comparing the c^{reg} -cost with the c -cost of paths and cycles in a path/cycle decomposition of x .

■ **Algorithm 2** Rounding (LP-ATSP $_{\rho}$) in regret metrics.

Input: asymmetric metric $(V \cup \{r\}, c^{\text{reg}})$ obtained from symmetric distances c .

Output: an Hamiltonian $s - t$ -rooted path P .

-
- R1.** Solve (LP-ATSP $_{\rho}$) to get an optimal extreme point solution x with value OPT_{LP} .
- R2.** Use Theorem 18 to find a convex combination of out-branchings B_1, \dots, B_q rooted at s and weights $\gamma_1, \dots, \gamma_q \geq 0$ summing to 1 such that t lies on each B_i and each $v \in V - \{s, t\}$ lies on at least a ρ -fraction of these branchings. Turn each B_i into a $s - t$ path P_i by adding the reverse (v, u) of each arc $(u, v) \in B_i$ that does not appear on the unique $s - t$ path in B_i and shortcutting the resulting Eulerian $s - t$ walk past repeated nodes.
- R3.** Define a cut requirement function $f : 2^V \rightarrow \{0, 1\}$ where $f(S) = 1$ if $\sum_{i: \text{red}(v, P_i) \subseteq S} \gamma_i < \delta$ for all $v \in S$. Observe f is downward-monotone: $f(S) \geq f(T)$ for sets $\emptyset \subsetneq S \subseteq T$. Use the LP-based 2-approximation in [16] to find a forest of undirected edges F such that $|\delta(S) \cap F| \geq f(S)$. Let \mathcal{C} be the components of F and let $C_1, \dots, C_{|\mathcal{C}|}$ be cycles on each component of F obtained by doubling and shortcutting each tree in F . For each cycle C_j of \mathcal{C} , let $w \in C_j$ be some *witness node* such that $\sum_{i: \text{red}(w, P_i) \subseteq V} \gamma_i \geq \delta$. Let W be the set of all witness over all C_j (note, it could be $W \cap \{s, t\} \neq \emptyset$). View each C_j as being traversed in some arbitrary direction.
- R4.** For each P_i , let P_i^W be the set of all nodes in $W \cap P_i$ such that all nodes of $\text{red}(w, P_i)$ are contained in the nodes of a single cycle C_j . Shortcut P_i past nodes not in $P_i^W \cup \{s, t\}$ and call this path P'_i . Note the nodes of P'_i lie in $W \cup \{s, t\}$.
- R5.** View P'_i with associated weights γ_i/δ as the path decomposition of an acyclic $s - t$ flow z with value $1/\delta$ with $z(\delta(w)) \geq 1$ for each $w \in W$. Further, $z(\delta^{\text{out}}(s)) = 1/\delta < 2$. By integrality of flows with upper- and lower-bounds on each node, we may decompose z as a convex combination of integral flows satisfying these bounds such that each flow supported consists of either 1 or 2 paths. Let P be the cheapest path among the flows with only one path in this decomposition. Note that P is an $s - t$ path spanning all of W .
- R6.** Complete P into a Hamiltonian $s - t$ path by adding all edges of the cycles C_i and shortcutting the resulting Eulerian walk.
-

Each P_i is obtained by adding the reverse of each edge uv of B_i not on the $s - t$ path in B_i (and then shortcutting the resulting Eulerian walk). Thus, $c(P_i) \leq 2 \cdot c(B_i) - c_{s,t}$ so $c^{\text{reg}}(P_i) \leq 2 \cdot (c(B_i) - c_{s,t})$. Thus, $\sum_i \gamma_i \cdot c^{\text{reg}}(P_i) \leq 2 \cdot \sum_i \gamma_i \cdot (c(B_i) - c_{s,t}) = (2 \cdot \sum_i \gamma_i \cdot c(B_i)) - 2 \cdot c_{s,t}$. Now, the convex combination of the B_i is dominated by x , so $\sum_i \gamma_i \cdot c(B_i) \leq \sum_e x_e \cdot c_e$. Finally, as x constitutes one unit of $s - t$ flow, the c -cost of x differs from the c^{reg} -cost of x exactly by $c_{s,t}$, so we finally see $\sum_i \gamma_i \cdot c^{\text{reg}}(P_i) \leq 2 \cdot OPT_{LP}$. ◀

The proofs of the following two lemmas proceed in a way that is very similar to related results [13] (though, their end goal was quite different). We defer their proofs to the end of this section.

► **Lemma 22.** *In Step 2, the function f is downward-monotone and $\sum_j c^{\text{reg}}(C_j) \leq \frac{6}{\rho - \delta} OPT_{LP}$.*

► **Lemma 23.** *The graph over V with edges $\cup_{i=1}^q P'_i$ is an acyclic graph. Further, for each $w \in W$ we have $\sum_{i: w \text{ lies on } P'_i} \gamma_i \geq \delta$. Finally, $\sum_{i=1}^q c^{\text{reg}}(P'_i) \leq 2 \cdot OPT_{LP}$.*

We now describe how to complete the analysis.

► **Lemma 24.** *In Step 5, the flow z has acyclic support, sends $1/\delta$ units of flow from s to t , and has $z(\delta^{\text{in}}(w)) \geq 1$ for each $w \in W$. The resulting path P has cost $\frac{2}{2\delta - 1} \cdot OPT_{LP}$.*

Proof. We have $\sum_i \gamma_i/\delta = 1/\delta$. As each P'_i is an $s - t$ flow, we have z given by $z_{uv} = \sum_{i: uv \in P'_i} \gamma_i/\delta$ is an $s - t$ flow of value $1/\delta$. Then by Lemma 23, the support of z is acyclic, $z(\delta^{\text{in}}(w)) \geq 1$ for each $w \in W$, and $\sum_{uv} c^{\text{reg}}_{u,v} z_{uv} \leq \frac{2}{\delta} \cdot OPT_{LP}$.

By integrality of flows with integral lower- and upper-bounds on the flow through each vertex, z may be decomposed into a convex-combination of integral flows f satisfying the lower-bound $f(\delta^{\text{in}}(w)) \geq 1$ for each $w \in W$ and $1 \leq f(\delta^{\text{out}}(s)) \leq 2$. Furthermore, the fraction of these flows f with $f(\delta^{\text{out}}(s)) = 1$ is exactly $2 - 1/\delta$, so the c^{reg} -cost of one such flow is at most $\frac{1}{2-1/\delta} \cdot \frac{2}{\delta} \cdot OPT_{LP} = \frac{2}{2\delta-1} \cdot OPT_{LP}$. Such a flow f has no cycles because the support of z is acyclic, so the edges supported by f form an $s - t$ path spanning all $w \in W$. ◀

The final path is formed from grafting the cycles $C_1, \dots, C_{|C|}$ into P , so the above results yield the following.

▶ **Theorem 25.** *The final path computed in Step 6 is a Hamiltonian $s - t$ path with c^{reg} -cost at most $\left(\frac{6}{\rho-\delta} + \frac{2}{2\delta-1}\right) \cdot OPT_{LP}$.*

Proof. By Lemma 24, the path P is an $s - t$ path spanning W with c^{reg} -cost at most $\frac{2}{2\delta-1} \cdot OPT_{LP}$. Each cycle C_j over a component in \mathcal{C} contains precisely one node in W , so the graph $P \cup_{j=1}^{|C|} C_j$ has an Eulerian $s - t$ walk that visits all nodes. By Lemma 22, the total c^{reg} -cost of all cycles is at most $\frac{6}{\rho-\delta} \cdot OPT_{LP}$. The result follows because shortcutting this Eulerian walk to get a Hamiltonian path does not increase the cost of the walk, by the triangle inequality. ◀

By setting $\delta = \frac{(2\sqrt{6}-1) \cdot \rho + 6 - \sqrt{6}}{10}$ (which optimizes the parameter), we get our main result showing the integrality gap is at most $\frac{300}{42-12\sqrt{6}} \cdot \frac{1}{2\rho-1} \approx \frac{23.8}{2\rho-1}$.

Proof of Lemma 22. That f is downward monotone is direct from the definition. We construct a vector x' over edges the undirected complete graph with nodes V with edge costs c . That is, for each undirected edge uv let $x'_{uv} = \frac{1}{\rho-\delta} \sum_{\substack{i:uv \text{ or } vu \\ \text{is red on } P_i}} \gamma_i$. We first claim $x'(\delta(S)) \geq f(S)$ for each $\emptyset \subsetneq S \subseteq V$. That is, suppose S is such that $f(S) = 1$ and let v satisfy $\sum_{i:\text{red}(v, P_i) \subseteq S} \gamma_i < \delta$. Since v lies on a ρ -fraction of paths in total, this means a $(\rho - \delta)$ -fraction of paths P_i have some edge of $\text{red}(v, P_i)$ crossing S , as required.

From Lemma 20, the total c -cost of all red edges on P_i is at most $\frac{3}{2} c^{\text{reg}}(P_i)$. Thus, $\sum_{uv} c_{uv} x'_{uv} \leq \frac{3}{2} \frac{1}{\rho-\delta} OPT_{LP}$. From using the LP-based 2-approximation in [16], the c -cost of the result forest is then at most $\frac{3}{\rho-\delta} OPT_{LP}$. By doubling the edges to get the cycles C_j , $\sum_j c(C_j) \leq \frac{6}{\rho-\delta} OPT_{LP}$. Finally, we chose an arbitrary direction for traversing each C_j but the c^{reg} -cost of a directed cycle is the same as its c -cost, so the result follows. ◀

Proof of Lemma 23. We claim that we do not keep two nodes from any red interval for each P_i when we form P'_i . But this is immediate from the fact that no cycle C_j contains two nodes of W .

By the definition of red intervals, any path P' obtained from a path P by shortcutting past all but one node in each red interval yields has its nodes appearing in strictly distance-increasing order. So, the P'_i paths all start at the same location, all end at the same location, and their internal nodes strictly increase in distance from s . So the union of all P'_i is an acyclic graph.

Now, consider some $w \in W$ and say it lies on cycle C_j . At least a δ -fraction of paths P_i spanning w satisfy $\text{red}(w, P_i) \subseteq C_j$ because $f(V(C_j)) = 0$, so each $w \in W$ lies on at least a δ -fraction of paths P'_i .

Since P'_i are obtained by shortcutting nodes from P_i , $\sum_{i=1}^q c^{\text{reg}}(P'_i) \leq \sum_{i=1}^q c^{\text{reg}}(P_i) \leq 2 \cdot OPT_{LP}$ by Lemma 21. ◀

References

- 1 F. Afrati, S. Cosmadakis, C. H. Papadimitriou, G. Papageorgiou, and N. Papakostantinou. The complexity of the traveling repairman problem. *Informatique Theorique et Applications*, 20(1):79–87, 1986.
- 2 H. C. An, R. Kleinberg, and D. B. Shmoys. Improving Christofides algorithm for the s-t path TSP. *Journal of the ACM*, 62(5):34, 2015.
- 3 A. Asadpour, M. X. Goemans, A. Madry, S. Oveis Gharan, and A. Saberi. An $O(\log n / \log \log n)$ -approximation algorithm for the asymmetric traveling salesman problem. In *In Proceedings of SODA*, pages 379–389, 2010.
- 4 J. Bang-Jensen, A. Frank, and B. Jackson. Preserving and increasing local edge-connectivity in mixed graphs. *SIAM J. Discrete Math.*, 8(2):155–178, 1995.
- 5 A. Blum, P. Chalasani, D. Coppersmith, W. R. Pulleyblank, P. Raghavan, and M. Sudan. The minimum latency problem. In *In Proceedings of STOC*, pages 163–171, 1994.
- 6 A. Blum, S. Chawla, D. R. Karger, T. Lane, and A. Meyerson. Approximation algorithms for orienteering and discount-reward TSP. *SIAM J. Comput.*, 37(2):653–670, 2007.
- 7 D. Chakrabarty and C. Swamy. Facility location with client latencies: linear-programming based techniques for minimum latency problems. *Math. of Operations Research*, 41(3):865–883, 2016.
- 8 M. Charikar, M. X. Goemans, and H. J. Karloff. On the integrality ratio for the asymmetric traveling salesman problem. *Math. of Operations Research*, 31(2):245–252, 2006.
- 9 K. Chaudhuri, B. Godfrey, S. Rao, and K. Talwar. Paths, trees, and minimum latency tours. In *In Proceedings of FOCS*, pages 36–45, 2003.
- 10 M. Fischetti, G. Laporte, and S. Martello. The delivery man problem and cumulative matroids. *Operations Research*, 41:1065–1064, 1993.
- 11 Z. Friggstad, A. Gupta, and M. Singh. An improved integrality gap for asymmetric TSP paths. *Math. of Operations Research*, 41(3):745–757, 2016.
- 12 Z. Friggstad, M. R. Salavatipour, and Z. Svitkina. Asymmetric traveling salesman path and directed latency problems. *SIAM J. Comput.*, 42(4):1596–1619, 2013.
- 13 Z. Friggstad and C. Swamy. Approximation algorithms for regret-bounded vehicle routing and applications to distance-constrained vehicle routing. In *In Proceedings of STOC*, pages 744–753, 2014.
- 14 Z. Friggstad and C. Swamy. Compact, provably-good LPs for orienteering and regret-bounded vehicle routing. In *In Proceedings of IPCO*, pages 199–211, 2017.
- 15 H. Gabow. Perfect arborescence packing in preflow mincut graphs. In *In Proceedings of SODA*, pages 528–538, 1996.
- 16 M. X. Goemans and D. P. Williamson. Approximating minimum-cost graph problems with spanning tree edges. *Operations Research Letters*, 16:183–189, 1994.
- 17 A. Köhne, V. Traub, , and J. Vygen. The asymmetric traveling salesman path LP has constant integrality ratio. In *In Proceedings of IPCO*, pages 288–298, 2019.
- 18 W. Mader. Konstruktion aller n-fach kantenzusammenhängenden Digraphen. *Europ. J. Combinatorics*, 3:63–67, 1982.
- 19 E. Minieka. The delivery man problem on a tree network. *Annals of Operations Res.*, 18:261–266, 1989.
- 20 V. Nagarajan and R. Ravi. The directed minimum latency problem. In *APPROX/RANDOM 2008. Proceedings*, Lecture Notes in Computer Science, pages 193–206. Springer, 2008.
- 21 J. Park and B. Kim. The school bus routing problem: A review. *European Journal of Operational Research*, 202(2):311–319, 2010.
- 22 I. Post and C. Swamy. Linear-programming based techniques for multi-vehicle minimum latency problems. In *Proceedings of SODA*, pages 512–531, 2015.
- 23 M. Skutella. List scheduling in order of α -points on a single machine. *Efficient Approximation and Online Algorithms*, 3484:250–291, 2006.

- 24 M. Spada, M. Bierlaire, and T. Liebling. Decision-aiding methodology for the school bus routing and scheduling problem. *Transportation Science*, 39(4):477–490, 2005.
- 25 O. Svensson, J. Tarnawski, and L. Vegh. A constant-factor approximation algorithm for the asymmetric traveling salesman problem. In *In Proceedings of STOC*, pages 204–213, 2018.
- 26 P. Toth and eds D. Vigo. *The Vehicle Routing Problem*. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, 2002.
- 27 V. Traub. *Approximation Algorithms for Traveling Salesman Problems*. PhD thesis, University of Bonn, 2020. URL: <http://hss.ulb.uni-bonn.de/2020/5834/5834.htm>.
- 28 V. Traub and J. Vygen. An improved approximation algorithm for ATSP. In *In Proceedings of STOC*, pages 1–13, 2020.

A Reduction to Instances with Polynomially-Bounded Integer Distances

► **Theorem 26.** *For any constant $\epsilon > 0$, if there is an $\alpha(n)$ -approximation for instances of DirLat where each $c_{u,v}$ is a positive integer bounded by a polynomial in n and $1/\epsilon$ and where $c_{u,v} \geq 1$ for nodes $u \neq v$, then there is an $(\alpha(n) + \epsilon)$ -approximation for general instances of DirLat.*

Proof. Compute a value ν such that $OPT \leq \nu \leq n^2 \cdot OPT$ where OPT is the optimum solution to the given DirLat instance. For example, ν could be the smallest value such that all nodes can be covered by a single walk in the graph $G_\nu = (V + r, E_\nu)$ consisting of directed edges $E_\nu = \{uv : c_{u,v} \leq \nu\}$. This can be checked, for example, by contracting the strongly-connected components of G_ν and checking if topologically sorting the resulting directed acyclic graph results in a single chain of components with the root in the first component.

Now, the case $OPT = 0$ can be detected in polynomial time as this is equivalent to checking if the strongly-connected components of the graph using only distance-0 edges forms a chain. So we assume $OPT > 0$, thus $\nu > 0$. We then assume $c_{u,v} \geq \epsilon \cdot \nu/n^3$ by increasing any distance that is smaller to this amount: the distances remain metric and the latency of any node on the optimum solution increases by at most $n \cdot \nu \leq \epsilon \cdot OPT/n$, so the total latency increases by at most $\epsilon \cdot OPT$.

Next, we may assume all distances satisfy $c_{u,v} \leq (\alpha(n) + 2\epsilon) \cdot \nu$ for the following reason. Suppose we update each distance $c_{u,v} > (\alpha(n) + 2\epsilon) \cdot \nu$ with $c_{u,v} = (\alpha(n) + 2\epsilon) \cdot \nu$. It is easy to check these updated distances also form a metric. The optimum solution cost is still OPT because no edge used by the optimum solution has its length shortened (as $\nu \geq OPT$). Also, note a solution P with $c(P) \leq (\alpha(n) + \epsilon) \cdot OPT$ will only use edges uv where $c_{u,v} < (\alpha(n) + 2\epsilon) \cdot \nu$. So an $(\alpha + \epsilon)$ -approximation in the metric with these truncated distances yields an $(\alpha + \epsilon)$ -approximation for the original distances.

Next, for all $u, v \in V + r$ let $d''(u, v) = \left\lfloor c_{u,v} \cdot \frac{n^4}{\nu \cdot \epsilon} \right\rfloor$. Let d' be the shortest path metric using edge distances given by d'' . Let OPT' denote the optimum solution to DirLat instance with distances d' . Observe

$$d'(u, v) \leq d''(u, v) \leq \frac{n^4}{\nu \cdot \epsilon} c_{u,v}.$$

Furthermore, $c_{u,v} \leq (\alpha(n) + 2\epsilon) \cdot \nu$ for each edge uv means $d'(u, v) \leq \frac{n^4}{\epsilon} \cdot (\alpha(n) + \epsilon)$. So all distances under d' are polynomially-bounded integers. We also see $OPT' \leq \frac{n^4}{\nu \cdot \epsilon} \cdot OPT$ by consider an optimum solution to the original instance, but under the new distances d' .

52:20 A Constant-Factor Approximation for Directed Latency in Quasi-Polynomial Time

Now consider a solution P with $d'(P) \leq \alpha(n) \cdot OPT'$. As d' is a metric, we may assume P is a Hamiltonian path so P traverses n edges. By replacing each edge in P with its shortest path using distances d'' , we obtain a walk W with $d''(W) = d'(P) \leq \alpha(n) \cdot OPT'$. For each edge uv , we have $d''(u, v) + 1 \geq c_{u,v} \cdot \frac{n^4}{\nu \cdot \epsilon}$. So the cost of W under d can be bounded as follows where sums over edges in W include as many terms of uv as its multiplicity in W .

$$\begin{aligned}
 c(W) &\leq \frac{\epsilon \cdot \nu}{n^4} \cdot \sum_{uv \in W} (d''(u, v) + 1) \\
 &= \frac{\epsilon \cdot \nu}{n^4} \cdot (d''(W) + |W|) \\
 &\leq \frac{\epsilon \cdot \nu}{n^4} \cdot (\alpha(n) \cdot OPT' + |W|) \\
 &\leq \alpha(n) \cdot OPT + \frac{\epsilon \cdot \nu}{n^2} \\
 &\leq (\alpha(n) + \epsilon) \cdot OPT.
 \end{aligned}$$

The last two bounds use $|W| \leq n \cdot |P| \leq n^2$ and $\nu \leq n^2 \cdot OPT$. ◀