Improved Guarantees for the A Priori TSP

Jannis Blauth
Research Inst. for Discrete Mathematics, Hausdorff Center for Math., University of Bonn, Germany

Meike Neuwohner
Research Inst. for Discrete Mathematics, Hausdorff Center for Math., University of Bonn, Germany

Luise Puhlmann
Research Inst. for Discrete Mathematics, Hausdorff Center for Math., University of Bonn, Germany

Jens Vygen
Research Inst. for Discrete Mathematics, Hausdorff Center for Math., University of Bonn, Germany

Abstract

We revisit the A PRIORI TSP (with independent activation) and prove stronger approximation guarantees than were previously known. In the A PRIORI TSP, we are given a metric space (V, c) and an activation probability $p(v)$ for each customer $v \in V$. We ask for a TSP tour T for V that minimizes the expected length after cutting T short by skipping the inactive customers.

All known approximation algorithms select a nonempty subset S of the customers and construct a *master route solution*, consisting of a TSP tour for S and two edges connecting every customer $v \in V \setminus S$ to a nearest customer in S.

We address the following questions. If we randomly sample the subset S, what should be the sampling probabilities? How much worse than the optimum can the best master route solution be? The answers to these questions (we provide almost matching lower and upper bounds) lead to improved approximation guarantees: less than 3.1 with randomized sampling, and less than 5.9 with a deterministic polynomial-time algorithm.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis

Keywords and phrases A priori TSP, random sampling, stochastic combinatorial optimization

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2023.14

Related Version *Full Version*: https://arxiv.org/abs/2309.10663

Supplementary Material *Software*: https://doi.org/10.60507/FK2/JCUIRI

1 Introduction

Many algorithms for stochastic discrete optimization problems sample a sub-instance, solve the resulting deterministic problem (often by some approximation algorithm), and extend this solution to the original instance [8, 11, 14, 15, 16, 25]. A nice and well-studied example is the A PRIORI TRAVELING SALESPERSON PROBLEM (A PRIORI TSP), which is the focus of this paper. What guarantee can we obtain by such an approach, even if we take an optimal sample? If we sample randomly, according to which distribution? What guarantee can we obtain by a deterministic polynomial-time algorithm? These are the questions addressed in this paper.

In the A PRIORI TSP (with independent activation), we are given a (semi-)metric space (V, c); the elements of V are called *customers*. Each customer v comes with an *activation probability* $0 < p(v) \leq 1$, so it will be active independently with probablity $p(v)$. However, we need to design a TSP tour T (visiting all of V) *before* knowing which customers will be active. *After* we know which customers are active we can cut the tour T short by skipping the inactive customers. The goal is to minimize the expected cost of the resulting tour (visiting the active customers).

34th International Symposium on Algorithms and Computation (ISAAC 2023).
Editors: Satoru Iwata and Naonori Kakimura; Article No. 14; pp. 14:1–14:16
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Note that computing an optimum a priori tour is APX-hard as the metric TSP is APX-hard [23], which is the special case where all activation probabilities are 1. We study approximation algorithms. A *ρ-approximation algorithm* for the A PRIORI TSP is a polynomial-time algorithm that computes a tour of expected cost at most $\rho \cdot \mathrm{OPT}$ for any given instance, where OPT denotes the expected cost of an optimum a priori tour.

Shmoys and Talwar [25] devised a randomized 4-approximation algorithm and a deterministic 8-approximation algorithm. A randomized constant-factor approximation algorithm was discovered independently by Garg, Gupta, Leonardi and Sankowski [11]. The randomized Shmoys–Talwar algorithm easily improves to a 3.5-approximation by using the Christofides–Serdyukov algorithm instead of the double tree algorithm as a subroutine for TSP (as noted by [7]), and slightly better using the new Karlin–Klein–Oveis Gharan algorithm [20]. The deterministic algorithm was improved to a 6.5-approximation by van Zuylen [28]; a slight improvement of this guarantee follows from the recent deterministic version of the Karlin–Klein–Oveis Gharan algorithm [21].

All known approximation algorithms for the A PRIORI TSP are of the following type. Select a nonempty subset S of customers and find a TSP tour for S (the *master tour*). Connect each other customer $v \in V \setminus S$ with a pair of parallel edges to a nearest point $\mu(v)$ in the master tour. We call this a *master route solution*. Once we know the set of active customers, we pay for the entire master tour (pretending to visit also its inactive customers!) and pay $2c(\mu(v), v)$ for each active customer v outside S to cover the round trip visiting v from $\mu(v)$. See Figure 1 for an example. Of course, we could cut the resulting tour shorter (we visit some inactive customers, and we visit some customers several times), but we will not account for this possible gain (unless fewer than two customers are active).

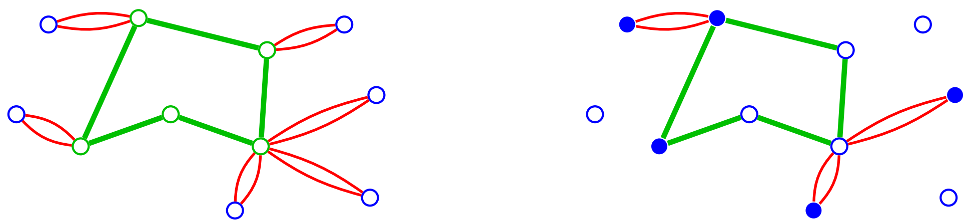

Figure 1 Left: A master route solution with a master tour (green, thick) and connections of the other customers to that master tour (red, curved). Right: After knowing which customers are active (filled), the master route solution reduces to a tour visiting all of the master tour and the other active customers.

1.1 Motivating questions

We start by reviewing the randomized algorithm by Shmoys and Talwar [25]. If fewer than two customers are active, any a priori tour can be cut short to a single point, resulting in cost zero. The algorithm by Shmoys and Talwar [25] selects each customer v independently into S with probability $p(v)$: exactly the activation probability. Assuming that the resulting set S is nonempty, there exists an associated master route solution with expected cost at most

$$\mathrm{MR}(S) \;:=\; \mathbb{E}_{A\sim p}\left[\mathbb{1}_{|A|\geq 2}\cdot\left(\mathrm{OPT}_{\mathrm{TSP}}(S,c)+2\cdot\sum_{v\in A}c(v,S)\right)\right].$$

Here $\mathrm{OPT}_{\mathrm{TSP}}(S,c)$ denotes the length of an optimum TSP tour for S, and $c(v,S) = \min\{c(v,s) : s \in S\}$ denotes the distance between v and a nearest customer in S (which is zero if $v \in S$); moreover, $\mathbb{E}_{A\sim p}$ denotes the expectation when the set A of active customers is sampled with respect to the given activation probabilities. Later on, $\mathbb{P}_{A\sim p}$ is used analogously. We multiply with $\mathbb{1}_{|A|\geq 2}$ (which is 1 if $|A| \geq 2$ and 0 otherwise) because the cost of the solution is zero if fewer than two customers are active.

If there is a customer d with $p(d) = 1$ (a *depot*), then S is never empty and we can bound

$$\mathrm{MR}(S) \;\leq\; \mathrm{OPT}_{\mathrm{TSP}}(S,c) + 2\cdot\mathbb{E}_{A\sim p}\left[\sum_{v\in A\setminus\{d\}} c(v, S\setminus\{v\})\right].$$

Note that the above upper bound also accounts for connecting active customers in S to the nearest other customer in S, which is not necessary but will allow the following. Taking the expectation over the random choice of S, an upper bound on the expected cost of that master route solution is

$$\mathbb{E}_{S\sim p}\left[\mathrm{MR}(S)\right] \;\leq\; \mathbb{E}_{S\sim p}\left[\mathrm{OPT}_{\mathrm{TSP}}(S,c)\right] + 2\cdot\mathbb{E}_{S\sim p}\left[\sum_{v\in S\setminus\{d\}} c(v, S\setminus\{v\})\right]$$

as the probability distributions to choose S and A are identical and the vertices are sampled independently. Since $\sum_{v\in S\setminus\{d\}} c(v, S\setminus\{v\}) \leq \mathrm{OPT}_{\mathrm{TSP}}(S,c)$ for all S and $\mathbb{E}_{S\sim p}\left[\mathrm{OPT}_{\mathrm{TSP}}(S,c)\right] \leq \mathrm{OPT}$, where OPT again denotes the expected cost of an optimum a priori tour, this yields

$$\mathbb{E}_{S\sim p}\left[\mathrm{MR}(S)\right] \;\leq\; 3\cdot\mathrm{OPT}. \tag{1}$$

The work of Shmoys and Talwar [25] implies that (1) also holds when there is no depot and when we take the conditional expectation under the condition that $|S| \geq 2$ (see also [28]). The Shmoys–Talwar algorithm cannot find an optimum TSP tour for S but uses the double tree algorithm with approximation guarantee 2. As noted by [7], one can as well use the Christofides–Serdyukov algorithm with approximation guarantee $\frac{3}{2}$, or in fact any α-approximation algorithm for TSP. Then the expected cost of the resulting master route solution is at most $(\alpha + 2)\cdot\mathrm{OPT}$.

This motivates the following questions:

(i) Is it optimal to sample S with exactly the activation probabilities (which is crucially used in the above analysis), or can we improve on the factor $\alpha + 2$ by sampling fewer or more?

(ii) How bad can the best master route solution be? We will call this the *master route ratio*: by the Shmoys–Talwar analysis, it is at most 3.

(iii) Can we obtain an approximation guarantee equal to the master route ratio by a master route solution based on random sampling, assuming that we can find optimum TSP tours? What is the best we can achieve with a $\frac{3}{2}$-approximation algorithm for TSP?

(iv) Can we obtain a better deterministic algorithm without a better TSP algorithm?

We give almost complete answers to all these questions.

1.2 Our results

The possibility that we sample the empty set or that no customer is active causes significant complications. The previous works [25] and [28] gave ad hoc proofs that *their* algorithms (which are also formulated with a depot) generalize to the non-depot case. We aim for a general reduction, losing only an arbitrarily small constant: Fortunately, instances in which the expected number of active customers is small can be solved easily with an approximation factor $3 + \varepsilon$ (for any $\varepsilon > 0$; similar to [8]), and hence much better than the known guarantees. For instances with a large expected number of active customers, one can assume without loss of generality (with an arbitrarily small loss) that there is a customer d that is always active, i.e., $p(d) = 1$ (see full version of this paper [4]). So we assume this henceforth and call d the depot. We summarize (and refer to the full version [4] for the proof):

▶ **Theorem 1.** *Let $\varepsilon > 0$ and $\rho \geq 3$ be constants. If there exists a (randomized) polynomial-time ρ-approximation algorithm for instances (V, c, p) of the A PRIORI TSP that have a depot (i.e., a customer d with $p(d) = 1$), then there is a (randomized) polynomial-time $(\rho + \varepsilon)$-approximation algorithm for general instances of the A PRIORI TSP.*

The Shmoys-Talwar algorithm [25] includes a customer v into S with probability $p(v)$: the sampling probability is exactly the activation probability. Although this is natural and allows for the simple analysis in Section 1.1 (assuming a depot), we show that this is not optimal. Decreasing the probability of including a customer into the master tour improves the approximation guarantee. To be more precise, in Section 2, we analyze the following *sampling algorithm* for A PRIORI TSP instances with depot. Let $f\colon (0, 1] \to [0, 1]$ with $f(1) = 1$.

(i) Sample a subset $S \subseteq V$ by including every customer v independently with probability $f(p(v))$.
(ii) Call an α-approximation algorithm for (metric) TSP in order to compute a TSP tour for S, which serves as master tour.
(iii) Connect every customer outside S to the nearest customer in S by a pair of parallel edges.

For a given instance this algorithm has expected approximation ratio at most

$$\frac{1}{\mathrm{OPT}} \cdot \mathbb{E}_{S \sim f \circ p} \left[\alpha \cdot \mathrm{OPT}_{\mathrm{TSP}}(S, c) + 2 \cdot \sum_{v \in V} p(v) \cdot c(v, S) \right] \tag{2}$$

(where $\frac{0}{0} := 1$). Shmoys and Talwar [25] used the identity function $f(p) = p$. It is easy to construct examples where sampling less or more is better. For example, if $c(v, w) = 1$ for all $v, w \in V$ with $v \neq w$ (and all activation probabilities except for the depot are tiny), it is best to include only the depot in the master tour: this yields an approximation ratio of 2 instead of 3. On the other hand, if $V = \{v_0, \ldots, v_{n-1}\}$ and $c(v_i, v_j) = \min\{j - i, n + i - j\}$ for $i < j$ (i.e., (V, c) is the metric closure of a cycle), the more we sample, the better. However, even if we choose f depending on the instance, there is a limit on what we can achieve:

▶ **Theorem 2.** *No matter how f is chosen, even depending on the instance in an arbitrary way, the sampling algorithm has no better approximation ratio than*

- 2.655 *even if it computes an optimum TSP tour on the sampled customers;*
- 3.049 *assuming that we never compute a TSP tour on the sampled customers of cost less than* 1.4999 *times the cost of an optimum tour.*

See the full version of our paper [4] for the proof. We do not have a matching upper bound, but we come close. For $\alpha = 1.5$ we prove (in Section 2):

▶ **Theorem 3.** *For $\alpha = 1.5$ and $f(p) = 1 - (1 - p)^\sigma$ with $\sigma = 0.663$, the sampling algorithm for A PRIORI TSP instances with depot has approximation guarantee less than 3.1.*

Figure 2 shows this function f. Together with Theorem 1 this immediately implies one of our main results:

▶ **Corollary 4.** *There is a randomized* 3.1*-approximation algorithm for A PRIORI TSP.* ◀

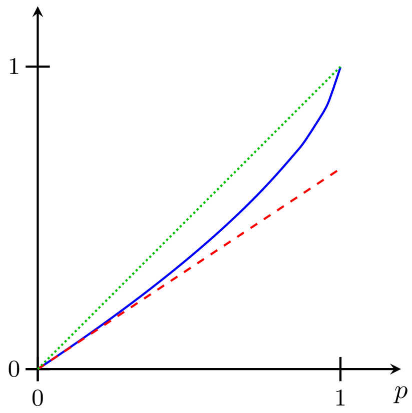

Figure 2 The function $p \mapsto 1-(1-p)^\sigma$ with $\sigma = 0.663$ (blue, solid) defines the sampling probability in Theorem 3, which is always at most the identity function (green, dotted), and for small p approximately equal to $p \mapsto \sigma \cdot p$ (red, dashed).

We conjecture that the bounds in Theorem 2 are actually attained by the sampling algorithm with $f(p) = 1-(1-p)^\sigma$, independent of the instance, where σ is a positive constant that depends on α only. See Comment 17 for details.

Having explored the limits of the random sampling approach, one might ask what is the limit of choosing an *optimal* master route solution. By van Zuylen's work [28], the answer to this question is the key to obtain a better deterministic approximation algorithm. Let us define:

▶ **Definition 5** (master route ratio). *The* master route ratio *is defined to be the supremum of*

$$\frac{\min\{\mathrm{MR}(S) : \emptyset \neq S \subseteq V\}}{\mathrm{OPT}}$$

taken over all A PRIORI *TSP instances (where* $\frac{0}{0} := 1$*).*

It is very easy to see that the master route ratio is at least 2 (for example, if $c(v,w)=1$ for all $v,w \in V$ with $v \neq w$). By the Shmoys–Talwar analysis, it is at most 3. We show in the full version of our paper [4]:

▶ **Theorem 6.** *The master route ratio for* A PRIORI *TSP instances with depot is at least* $\frac{1}{1-e^{-1/2}} > 2.541$ *and less than* 2.6*.*

We conjecture that the master route ratio is exactly $\frac{1}{1-e^{-1/2}}$.

As van Zuylen's [28] analysis reveals (cf. [4]), her algorithm is a $(2+\alpha\rho)$-approximation algorithm if the master route ratio is ρ and we have an algorithm for TSP that guarantees to produce a tour of cost at most α times the value of the subtour relaxation. So our new upper bound on the master route ratio immediately implies a better guarantee (combining Theorems 1 and 6 with $\alpha = \frac{3}{2}$ [27]):

▶ **Corollary 7.** *There is a deterministic 5.9-approximation algorithm for* A PRIORI *TSP.* ⌟

1.3 Our techniques

The lower bounds (Theorem 2 and the lower bound in Theorem 6) are obtained by analyzing simple examples. The main technical difficulty is in proving the upper bounds.

To prove Theorem 3 and the upper bound in Theorem 6, we will show that it suffices to consider instances in which all customers (except the depot) have the same tiny activation probability. We call these instances *normalized*.

▶ **Definition 8.** *Let $\varepsilon > 0$. An instance (V, c, p) of* A PRIORI *TSP is called ε-*normalized *if the instance contains a depot $d \in V$ (with $p(d) = 1$), and $p(v) = \varepsilon$ for all $v \in V \setminus \{d\}$.*

Given an instance of A PRIORI TSP with a depot d, one can transform it to a normalized instance by replacing each customer $v \in V \setminus \{d\}$ by many copies, each with the same tiny activation probability, such that the probability that at least one of these copies is active is roughly $p(v)$. This way, the master route ratio and the approximation guarantee of the sampling algorithm can only get worse. More precisely, we show in the full version of this paper [4]:

▶ **Lemma 9.** *Let $(\varepsilon_i)_{i\in\mathbb{N}} \in (0,1]^{\mathbb{N}}$ with $\lim_{i\to\infty} \varepsilon_i = 0$. Let $\mathcal{I}$ be the class of all ε-normalized instances with $\varepsilon = \varepsilon_i$ for some $i \in \mathbb{N}$. Then*

(i) *The master route ratio is the same when restricting it to instances in $\mathcal{I}$ and when restricting it to all instances with depot.*

(ii) *Let $\sigma \in (0,1)$. Every upper bound on (2) for $f(p) = \sigma p$ $\forall p \in (0,1)$ for all instances in $\mathcal{I}$ implies the same upper bound on (2) for $f(p) = 1 - (1-p)^\sigma$ for arbitrary instances with depot.*

On a high level, our proofs of Theorem 3 and Theorem 6 are similar. In both cases we will design a linear program that encodes the metric c by variables and minimizes the expected cost of an optimum a priori tour subject to (a relaxation of) the constraint that the expected cost of the output of the sampling algorithm is at least 1 (for Theorem 3) or the expected cost of any master route solution is at least 1 (for Theorem 6), respectively. Then the reciprocals of the LP values yield the desired upper bounds.

However, this approach has to overcome several obstacles. First, it is not obvious how to encode the metric c by finitely many variables, given that we need to consider arbitrary instance sizes. We do this by fixing an optimum a priori tour T^* (a cyclic order of the customers) and carefully aggregating distances of customer pairs with the same number of hops in between on T^*. Of course we exploit the structure of normalized instances.

In the end, we will (almost) ignore variables that correspond to a very large number of hops (where it is very unlikely that none of the customers "in between" is active). These variables have negligible impact because the probability that these edges occur decreases exponentially with increasing number of hops on T^*, whereas the average length of these edges can only grow linearly due to the triangle inequality.

The next idea is to consider certain structured solutions only. Rather than connecting a customer v that is not in the master tour to the *nearest* customer $\mu(v)$ in the master tour, we consider only two possible members of the master tour: we traverse T^* from v in each of the two possible directions, and consider the first customer that we meet and that is contained in our master tour. None of these two may be a nearest one in the master tour, but we still obtain an upper bound. For bounding the master route ratio, we will in addition only consider master tours whose customers are equidistantly distributed on T^* (except for the depot).

In this way, we obtain an optimization problem for a fixed uniform activation probability p (i.e., for p-normalized instances). However, we must let $p \to 0$ according to Lemma 9 and hence need a description that is independent of p. This is another major obstacle. To overcome it, we use a second level of aggregation (buckets, rounding the number of hops to integer multiples of, say, $\frac{1}{100p}$). However, this causes several difficulties. In the case of the sampling algorithm, describing the expected cost of the output of the sampling algorithm in terms of the buckets is nontrivial. In case of the master route ratio, the same holds for master route solutions and actually requires a third level of aggregation (bucket intervals).

In the end, we obtain (in both cases) a single, relatively compact, linear program that yields an upper bound for all instance sizes and all activation probabilities from a sequence that converges to zero. We solve the dual LP numerically and just need to check feasibility to prove the desired upper bounds.

1.4 Further related work

The TSP has also been studied under the aspect of robust optimization, where the set of customers that need to be visited is known in advance, but the edge lengths are chosen probabilistically or even adversarially [10, 26]. The a priori optimization problem where the set of customers is chosen adversarially is known as universal TSP [11, 13, 24]. The probabilistic version that we consider was introduced by Jaillet [17] and Bertsimas [2]. Since then, various aspects of the problem have been investigated, including the asymptotic behavior of random instances [2, 3, 5, 17, 18], online variants [11], or exact algorithms [1]. Approximation algorithms have also been studied for general probability distributions [6, 13, 24].

Other problems that have been considered in an a priori setting include vehicle routing, traveling repairman, Steiner tree, and network design [7, 8, 9, 11, 14, 15, 16, 22]. However, none of these works managed to determine the approximation guarantee of their algorithms exactly.

Previous approaches to design a linear program that yields the approximation ratio of a certain algorithm for some optimization problem (e.g., [12, 19]) typically required an infinite family of linear programs and could not obtain a bound for general instances by just solving a single linear program.

2 Upper bound on the approximation ratio of random sampling

In this section we will prove Theorem 3. As mentioned earlier, we will design a single linear program such that the reciprocal of its optimum value is an upper bound on the approximation ratio of the sampling algorithm for a certain class of normalized instances. For this sake, let $\beta, b_0 > 0$ be constants that we will choose later. We will consider ε-normalized instances where ε is of the form $\varepsilon = \frac{\beta}{b}$ for some odd integer $b \geq b_0$. The meaning of these constants will become clear in Section 2.2. For such instances we will obtain an upper bound on the approximation ratio of the sampling algorithm, when sampling each customer with probability σp for $\sigma = 0.663$ (in addition to the depot). Combined with Lemma 9, this immediately yields the same upper bound on the approximation guarantee of the sampling algorithm that samples each customer v with probability $1 - (1 - p(v))^{\sigma}$ for arbitrary A PRIORI TSP instances with depot.

2.1 An optimization problem to bound the approximation ratio

In this section, we first describe an upper bound for all p-normalized instances (for a fixed uniform activation probability p) by a single optimization problem. We will consider the algorithm that samples each customer with probability σp. Let T^* be a fixed optimum a priori tour, with customers appearing in the order $v_0, v_1, \ldots, v_{n-1}$; here v_0 denotes the depot. Let $v_i := v_0$ for $i < 0$ or $i > n - 1$. For $k \in \mathbb{Z}_{\geq 1}$ we define

$$C_k \;:=\; p^2 \cdot \sum_{j \in \mathbb{Z}} c(v_j, v_{j+k}).$$

Observe that only finitely many summands are nonzero. See Figure 3 for an example. Since c is a metric, the numbers C_k are nonnegative and satisfy the triangle inequality, that is, for all $i, j \geq 1$

$$C_{i+j} \;\leq\; C_i + C_j. \tag{3}$$

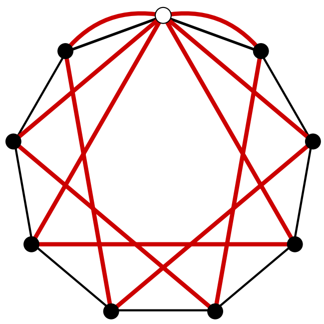

Figure 3 The depot v_0 is the white circle at the top; the tour T^* is drawn in black. Adding up the costs of the edges marked in red gives $\frac{C_3}{p^2}$.

Moreover, we can express the expected cost of T^* in terms of the C_i.

▶ **Proposition 10.** *The expected cost of T^* is exactly*

$$\sum_{i=1}^{\infty}(1-p)^{i-1} \cdot C_i. \tag{4}$$

Proof. Let $1 \leq i \leq n-2$ and $1 \leq j \leq n-i-1$. Then v_j and v_{j+i} are consecutive active customers with probability $p^2 \cdot (1-p)^{i-1}$; note that the cost of the edge $\{v_j, v_{j+i}\}$ is counted with exactly the same coefficient in (4). Moreover, for $1 \leq j \leq n-1$, v_j is the first active customer after the depot with probability $p \cdot (1-p)^{j-1}$, and the cost of the edge $\{v_0, v_j\}$ is counted $\sum_{i=j}^{\infty} p^2 \cdot (1-p)^{i-1} = p \cdot (1-p)^{j-1}$ times in (4). By symmetry, the terms also match for the last active customer before the depot. ◀

We now consider the master route solution resulting from sampling each customer with probability σp (in addition to the depot). Let α again denote the approximation guarantee of the TSP algorithm that we use. We will now show that the expected cost of this master route solution is at most

$$\sigma^2 \sum_{k=1}^{\infty}(1-\sigma p)^{k-1} \cdot \left(\alpha \cdot C_k + 2p \cdot \sum_{i=1}^{k-1} \min\left\{C_i,\, C_{k-i}\right\}\right). \tag{5}$$

By the same argumentation as in the proof of Proposition 10, the master tour has expected cost at most

$$\alpha \cdot \mathbb{E}_{S\sim q}[c(T^*[S])] \;=\; \alpha \cdot \sigma^2 \cdot \sum_{k=1}^{\infty}(1-\sigma p)^{k-1} \cdot C_k,$$

where $q(v) = \sigma p$ for all $v \in V \setminus \{d\}$ and $q(d) = 1$.

Next we bound the expected cost of connecting the active customers to the master tour. Instead of connecting v to the nearest customer in the master tour, we consider only two options: the first sampled customer that we meet when traversing T^* from v in either direction. Note that sampling v_0 with probability 1 is equivalent to sampling each v_j with $j \leq 0$ and $j \geq n$ with probability σp. Now, for $j \in \mathbb{Z}$ and $k \geq 2$, the probability that v_j and v_{j+k} are sampled, but none of the intermediate customers is, equals $(\sigma p)^2 \cdot (1-\sigma p)^{k-1}$. In this case, the total expected cost of connecting the intermediate active customers can be bounded by $2p \cdot \sum_{i=1}^{k-1} \min\{c(v_j, v_{j+i}), c(v_{j+i}, v_{j+k})\}$. Thus we can bound the expected cost of connecting all active customers to the master tour by

$$\begin{aligned}
&\sum_{k=2}^{\infty} \sigma^2 p^2 \cdot (1-\sigma p)^{k-1} \cdot \sum_{j\in\mathbb{Z}} 2p \cdot \sum_{i=1}^{k-1} \cdot \min\left\{c(v_j, v_{j+i}),\, c(v_{j+i}, v_{j+k})\right\} \\
\leq\ & 2\sigma^2 p^3 \cdot \sum_{k=1}^{\infty} (1-\sigma p)^{k-1} \cdot \sum_{i=1}^{k-1} \min\left\{\sum_{j\in\mathbb{Z}} c(v_j, v_{j+i}),\, \sum_{j\in\mathbb{Z}} c(v_{j+i}, v_{j+k})\right\} \\
=\ & 2\sigma^2 p^3 \cdot \sum_{k=1}^{\infty} (1-\sigma p)^{k-1} \cdot \sum_{i=1}^{k-1} \min\left\{\sum_{j\in\mathbb{Z}} c(v_j, v_{j+i}),\, \sum_{j\in\mathbb{Z}} c(v_j, v_{j+(k-i)})\right\} \\
=\ & 2\sigma^2 p \cdot \sum_{k=1}^{\infty} (1-\sigma p)^{k-1} \cdot \sum_{i=1}^{k-1} \min\left\{C_i,\, C_{k-i}\right\}.
\end{aligned}$$

We conclude that the ratio of (5) to (4) is an upper bound on the approximation guarantee of the sampling algorithm for that instance. Note that the number of customers appears neither in (4) nor in (5). In other words, minimizing (4) subject to the constraints that (5) is equal to 1 and the C_i are nonnegative and satisfy the triangle inequality (3) yields the reciprocal of an upper bound on the approximation guarantee of the sampling algorithm on all p-normalized instances. We arrive at the following optimization problem:

$$\min \sum_{i=1}^{\infty} (1-p)^{i-1} \cdot C_i \qquad \text{(Sampling-OP)}$$

$$\begin{aligned}
\text{subject to} \qquad C_i &\geq 0 && \text{for } i \in \mathbb{N} && (6)\\
C_i + C_j &\geq C_{i+j} && \text{for } i,j \in \mathbb{N} && (7)\\
\sum_{k=1}^{\infty} (1-\sigma p)^{k-1} \cdot \left(\alpha \cdot C_k + 2p \cdot \sum_{i=1}^{k-1} \min\left\{C_i,\, C_{k-i}\right\}\right) &\geq \sigma^{-2}. && && (8)
\end{aligned}$$

Note that in (8) we only require that (5) is at least 1 instead of exactly 1. This does not change the infimum because we can always scale all the C_i's. We have proved:

▶ **Lemma 11.** *Let $0 < p < 1$. The reciprocal of the value of* (Sampling-OP) *is an upper bound on the approximation guarantee for the sampling algorithm with $f(p) = \sigma p$ for all p-normalized instances.* ◀

2.2 Obtaining a single linear program

Note that we have an infinite set of optimization problems (one for each choice of p), and, in view of Lemma 9, we have to consider the limit for $p \to 0$.

In the following, we require that p is of the form $p = \frac{\beta}{b}$ for some odd integer $b \geq b_0$. Note that $p \to 0$ as $b \to \infty$. In order to obtain a single optimization problem for all such values of p, we put subsequent C_i's into buckets of size b. More precisely, we define buckets

$$B_i := \sum_{j=\max\{1, ib-\frac{b-1}{2}\}}^{ib+\frac{b-1}{2}} C_j \qquad (9)$$

for $i \geq 0$. In the following, we show that we can use the constraints in (Sampling-OP) to generate (slightly relaxed) constraints that only depend on these buckets. First, we note that the buckets are chosen such that they still satisfy the triangle inequality.

▶ **Proposition 12.** *For all $i, j \geq 1$,*

$$B_{i+j} \;\leq\; B_i + B_j.$$

Proof. Indeed, using (7), as illustrated in Figure 4,

$$\begin{aligned}
B_{i+j} \;&=\; \sum_{k=-\frac{b-1}{2}}^{\frac{b-1}{2}} C_{(i+j)b+k} \;=\; \sum_{k=0}^{\frac{b-1}{2}} C_{(i+j)b-\frac{b-1}{2}+2k} + \sum_{k=1}^{\frac{b-1}{2}} C_{(i+j)b-\frac{b-1}{2}+2k-1} \\
&\leq \sum_{k=0}^{\frac{b-1}{2}} \left(C_{ib-\frac{b-1}{2}+k} + C_{jb+k}\right) + \sum_{k=1}^{\frac{b-1}{2}} \left(C_{ib+k} + C_{jb-\frac{b-1}{2}+k-1}\right) \\
&= \sum_{k=-\frac{b-1}{2}}^{\frac{b-1}{2}} C_{ib+k} + \sum_{k=-\frac{b-1}{2}}^{\frac{b-1}{2}} C_{jb+k} \;=\; B_i + B_j.
\end{aligned}$$

◀

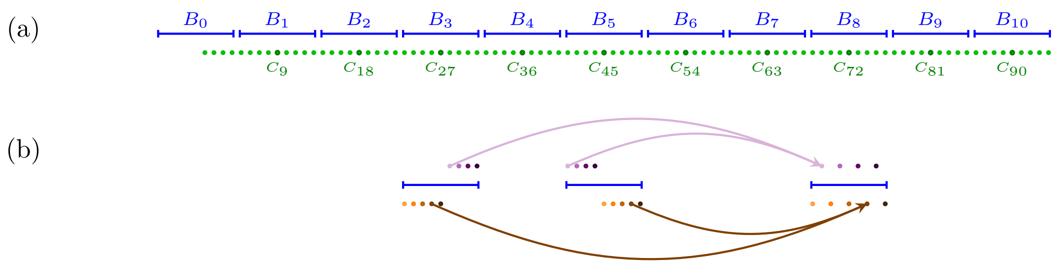

Figure 4 (a): The green dots stand for $C_1, C_2, \ldots$, and the centers of the buckets (C_{ib} for $i \geq 1$) are highlighted. Here the bucket size is $b = 9$, and the blue intervals show the buckets $B_0, B_1, B_2, \ldots$ (b): Combining the triangle inequalities for the C_i's leads to triangle inequalities for the B_i's; here shown for $B_i = B_3$ and $B_j = B_5$: We add up all triangle inequalities for C_k from B_3 and C_ℓ from B_5 where C_k and C_ℓ have the same color; illustrated with $C_{26} + C_{48} \leq C_{74}$ and $C_{28} + C_{41} \leq C_{69}$.

Next we aim for an upper bound on the left-hand side of (8) that only depends on the buckets. First we show:

▶ **Lemma 13.**

$$\sum_{k=1}^{\infty} (1-\sigma p)^{k-1} \cdot \sum_{i=1}^{k-1} \min\left\{C_i,\, C_{k-i}\right\} \;\leq\; b \cdot \sum_{i=0}^{\infty}\sum_{j=0}^{\infty} e^{-(i+j-1)\cdot\sigma b p} \cdot \min\{B_i, B_j\}.$$

Proof. For $i \in \mathbb{Z}_{\geq 0}$, let $I_i = \{\max\{1, ib - \frac{b-1}{2}\}, \ldots, ib + \frac{b-1}{2}\}$ be the set of indices in the i-th bucket. Then

$$\begin{aligned}
&\sum_{k=1}^{\infty} (1-\sigma p)^{k-1} \cdot \sum_{i=1}^{k-1} \min\left\{C_i,\, C_{k-i}\right\} \\
&= \sum_{k=1}^{\infty}\sum_{\ell=1}^{\infty} (1-\sigma p)^{k+\ell-1} \cdot \min\left\{C_k,\, C_\ell\right\}
\end{aligned}$$

$$\leq \sum_{i=0}^{\infty}\sum_{j=0}^{\infty}(1-\sigma p)^{\max\{0,i+j-1\}\cdot b} \cdot \sum_{k\in I_i}\sum_{\ell\in I_j}\min\{C_k, C_\ell\}$$
$$\leq \sum_{i=0}^{\infty}\sum_{j=0}^{\infty} e^{-(i+j-1)\cdot\sigma b p} \cdot \sum_{k\in I_i}\sum_{\ell\in I_j}\min\{C_k, C_\ell\}.$$

In the last inequality we used $1-x \leq e^{-x}$ for all $x \in \mathbb{R}$. Now, for $i, j \in \mathbb{Z}_{\geq 0}$, consider the complete bipartite graph H where one bipartition consists of $|I_j|$ copies of every element of I_i, and the other bipartition consists of $|I_i|$ copies of every element of I_j. Then

$$\sum_{(k,\ell)\in E(H)} \min\{C_k, C_\ell\} \;=\; |I_i|\cdot|I_j|\cdot\sum_{k\in I_i}\sum_{\ell\in I_j}\min\{C_k, C_\ell\}.$$

We can partition $E(H)$ into $t := |I_i|\cdot|I_j|$ perfect matchings $M_1, \ldots, M_t$. Then

$$\begin{aligned}\sum_{(k,\ell)\in E(H)} \min\{C_k, C_\ell\} &= \sum_{s=1}^{t}\sum_{(k,\ell)\in M_s}\min\{C_k, C_\ell\}\\ &\leq \sum_{s=1}^{t}\min\left\{\sum_{(k,\ell)\in M_s} C_k, \sum_{(k,\ell)\in M_s} C_\ell\right\}\\ &= \sum_{s=1}^{t}\min\left\{|I_j|\cdot\sum_{k\in I_i} C_k, |I_i|\cdot\sum_{\ell\in I_j} C_\ell\right\}\\ &= |I_i|\cdot|I_j|\cdot\min\left\{|I_j|\cdot B_i, |I_i|\cdot B_j\right\}.\end{aligned}$$

Note that the second equality follows from the fact that $V(H)$ contains $|I_j|$ copies of each element in I_i and vice versa. Moreover, summing over the endpoints of the edges in a perfect matching in M is the same as summing over $V(H)$. Division by $|I_i|\cdot|I_j|$ yields

$$\sum_{k\in I_i}\sum_{\ell\in I_j}\min\{C_k, C_\ell\} \;\leq\; \min\left\{|I_j|\cdot B_i, |I_i|\cdot B_j\right\} \;\leq\; b\cdot\min\{B_i, B_j\}. \qquad ◀$$

Using Lemma 13 and $\beta = bp$, the left-hand side of (8) can be upper bounded by

$$\begin{aligned}&\alpha\cdot\sum_{k=1}^{\infty}(1-\sigma p)^{k-1}\cdot C_k \;+\; 2bp\cdot\sum_{i=0}^{\infty}\sum_{j=0}^{\infty}e^{-(i+j-1)\cdot\sigma b p}\cdot\min\{B_i, B_j\}\\ \leq\;& \alpha\cdot\sum_{k=0}^{\infty}(1-\sigma p)^{\max\{0,kb-\frac{b-1}{2}-1\}}\cdot B_k \;+\; 2bp\cdot\sum_{i=0}^{\infty}\sum_{j=0}^{\infty}e^{-(i+j-1)\cdot\sigma b p}\cdot\min\{B_i, B_j\}\\ \leq\;& \alpha\cdot\sum_{k=0}^{\infty}e^{-(k-1)\cdot\sigma\beta}\cdot B_k \;+\; 2\beta\cdot\sum_{i=0}^{\infty}\sum_{j=0}^{\infty}e^{-(i+j-1)\cdot\sigma\beta}\cdot\min\{B_i, B_j\}.\end{aligned} \tag{10}$$

The last inequality follows from $(1-\sigma p)^{kb-\frac{b-1}{2}-1} \leq (1-\sigma p)^{kb-b}$ and $1+x \leq e^x$ for all $x \in \mathbb{R}$. Note that we still sum over infinitely many variables. Hence, in order to get a finite linear program, we aim for an upper bound on (10) that only depends on the buckets B_i with $i \leq N$ for some integer N that we will choose later. For this, we use the triangle inequality (Proposition 12) to bound the terms depending on buckets B_i with $i > N$ by some term depending on $B_1, \ldots, B_N$ only. For large N this will result in a negligible error as the coefficients in (10) decrease exponentially. In Section 2.4 we prove the following bound on the error term:

▶ **Lemma 14.** *Let* $\delta_1 := \frac{4\beta}{e^{N\sigma\beta}(e^{\sigma\beta}-1)}$ *and* $\delta_2 := \left(\alpha + \frac{2\beta}{e^{N\sigma\beta}(e^{\sigma\beta}-1)}\right) \cdot \frac{e^{-N\sigma\beta}}{(1-e^{-\sigma\beta})^2} \cdot (1+N-e^{-\sigma\beta}N)$. *Then*

$$\alpha \cdot \sum_{k=N+1}^{\infty} e^{-(k-1)\cdot\sigma\beta} \cdot B_k \;+\; 2\beta \cdot \sum_{i,j \in \mathbb{Z}_{\geq 0}: \max\{i,j\} > N} e^{-(i+j-1)\cdot\sigma\beta} \cdot \min\{B_i, B_j\}$$
$$\leq\; \delta_1 \cdot \sum_{k=0}^{N} e^{-(k-1)\cdot\sigma\beta} \cdot B_i + \delta_2 \cdot B_1.$$

Therefore, we get a lower bound on (Sampling-OP) by minimizing $\sum_{i=1}^{\infty}(1-p)^{i-1} \cdot C_i$ subject to (9) and $B_i \geq 0$ for $i \geq 0$, $B_{i+j} \leq B_i + B_j$ for $i, j \geq 1$ with $i + j \leq N$, and

$$(\alpha + \delta_1) \cdot \sum_{k=0}^{N} e^{-(k-1)\cdot\sigma\beta} \cdot B_k + 4\beta \cdot \sum_{j=0}^{N}\sum_{i=0}^{j} e^{-(i+j-1)\cdot\sigma\beta} \cdot \min\{B_i, B_j\} + \delta_2 \cdot B_1 \geq \sigma^{-2}. \quad (11)$$

Note that the objective still contains infinitely many variables and depends on p. The first problem can easily be resolved by bounding

$$\sum_{i=1}^{\infty}(1-p)^{i-1} \cdot C_i \;\geq\; \sum_{i=0}^{\infty}(1-p)^{bi+\frac{b-1}{2}-1} \cdot B_i \;\geq\; \sum_{i=0}^{N}(1-p)^{(i+\frac{1}{2})b} \cdot B_i. \quad (12)$$

It remains to get rid of the dependence on b and p (recall that $p = \frac{\beta}{b}$). To this end, we exploit that $\lim_{b\to\infty}(1-\frac{\beta}{b})^{(i+\frac{1}{2})b} = e^{-(i+\frac{1}{2})\beta}$ for all $i = 0, \ldots, N$, and that by Lemma 9, we can choose b_0 arbitrarily large. This will allow us to conclude that we can replace the objective by $\sum_{i=0}^{N} e^{-(i+\frac{1}{2})\beta} \cdot B_i$ and still obtain an upper bound (see the proof of Lemma 15 for the technical details). Putting everything together, we arrive at the following LP.

$$\min \sum_{i=0}^{N} e^{-(i+\frac{1}{2})\beta} \cdot B_i \qquad \text{(Sampling-LP)}$$

$$\text{subject to} \quad (\alpha + \delta_1) \cdot \sum_{k=0}^{N} e^{-(k-1)\cdot\sigma\beta} \cdot B_k + \delta_2 \cdot B_1 + 4\beta \cdot \sum_{j=0}^{N}\sum_{i=0}^{j} e^{-(i+j-1)\cdot\sigma\beta} \cdot M_{i,j} \;\geq\; \sigma^{-2} \quad (13)$$

$$B_i + B_j \;\geq\; B_{i+j} \qquad \text{for } 1 \leq i \leq j \leq N \text{ with } i + j \leq N \quad (14)$$
$$B_i \;\geq\; M_{i,j} \qquad \text{for } 0 \leq i \leq j \leq N \quad (15)$$
$$B_j \;\geq\; M_{i,j} \qquad \text{for } 0 \leq i \leq j \leq N \quad (16)$$
$$B, M \;\geq\; 0. \quad (17)$$

Recall that δ_1 and δ_2 were defined in Lemma 14. We conclude:

▶ **Lemma 15.** *Let N be an integer and $\beta > 0$. The reciprocal of the optimum value of* (Sampling-LP) *is an upper bound on the approximation guarantee of the sampling algorithm for $f(p) = 1 - (1-p)^\sigma$ (using an α-approximation algorithm for TSP), for all* A PRIORI *TSP instances with depot.*

Proof. We compare the value of (Sampling-LP) to the value of (Sampling-OP). We showed above that for any feasible solution C to (Sampling-OP) we obtain a feasible solution (B, M) to (Sampling-LP) via (9) and $M_{i,j} = \min\{B_i, B_j\}$.

Fix $\delta > 0$. Then there exists $b_0 \in \mathbb{N}$ such that for all odd integers $b \geq b_0$

$$\sum_{i=0}^{N} e^{-(i+\frac{1}{2})\beta} \cdot B_i \;\leq\; (1+\delta) \cdot \sum_{i=0}^{N}\left(1 - \frac{\beta}{b}\right)^{(i+\frac{1}{2})b} \cdot B_i \;\overset{(12)}{\leq}\; (1+\delta) \cdot \sum_{i=0}^{\infty}\left(1 - \frac{\beta}{b}\right)^{i-1} \cdot C_i.$$

Thus the value of (Sampling-LP) is at most $(1+\delta)$ times the value of (Sampling-OP) with $p = \frac{\beta}{b}$ for all odd integers $b \geq b_0$. Hence, by Lemma 11, $(1+\delta)$ times the reciprocal of the optimum value of (Sampling-LP) is an upper bound on (2) for all $\frac{\beta}{b}$-normalized instances for all odd integers $b \geq b_0$. By Lemma 9, the same bound then holds for all instances with depot. Since this bound holds for all $\delta > 0$, it also holds for $\delta = 0$. ◀

2.3 The dual LP

In order to obtain a lower bound on the optimum value of (Sampling-LP), we provide a *feasible solution* to the *dual* linear program. For the dual LP, we introduce variables $x_{i,j}$ for the inequalities of type (14), variables $v_{i,j}$ and $w_{i,j}$ for the inequalities of type (15) and (16), respectively, and a variable y for inequality (13). Using these variables, the dual LP looks as follows:

$$\max\ \sigma^{-2} \cdot y \qquad \text{(Dual-Sampling-LP)}$$

$$\text{subject to} \quad 4\beta \cdot e^{-(i+j-1)\cdot\sigma\beta} \cdot y \ \leq\ v_{i,j} + w_{i,j} \qquad \text{for } 0 \leq i \leq j \leq N \quad (18)$$

$$(\alpha+\delta_1)\cdot e^{-(k-1)\cdot\sigma\beta}\cdot y + \sum_{j=k}^{N} v_{k,j} + \sum_{j=0}^{k} w_{j,k} + \mathbb{1}_{k=1}\cdot\delta_2\cdot y$$
$$+\mathbb{1}_{k>0}\cdot\left(\sum_{i=1}^{\min\{k,N-k\}} x_{i,k} + \sum_{i=k}^{N-k} x_{k,i} - \sum_{\substack{1\leq i\leq j\leq N,\\ i+j=k}} x_{i,j}\right) \ \leq\ e^{-(k+\frac{1}{2})\beta} \qquad \text{for } 0 \leq k \leq N \quad (19)$$

$$x, y, v, w \ \geq\ 0. \qquad (20)$$

▶ **Corollary 16.** *Let N be an integer and $\beta > 0$. For any feasible solution (x, y, v, w) to* (Dual-Sampling-LP)*, σ^2/y is an upper bound on the approximation ratio of the sampling algorithm with $f(p) = 1-(1-p)^\sigma$ restricted to* a priori *TSP instances with depot.* ⌟

We have computed a dual solution using Gurobi 10.0.1 with $\alpha = 1.5$, $\beta = \frac{1}{100}$, $N = 2500$, and $\sigma = 0.663$, yielding an upper bound of 3.094 and thus proving Theorem 3. The dual solution and a Python script that verifies that this is a feasible solution to (Dual-Sampling-LP) can be found at `https://doi.org/10.60507/FK2/JCUIRI`. For $\alpha = 1$, we get an upper bound of 2.694.

▶ **Comment 17.** Solving (Sampling-LP) with the same values for α, β, N, and σ yields an a priori TSP instance of the same shape as the lower bound example provided in [4]. Hence we conjecture that the upper bound given by (Dual-Sampling-LP) converges to the lower bound given in Theorem 2 for $\beta \to 0$ and $N\beta \to \infty$.

2.4 Bounding the error term (Proof of Lemma 14)

We first prove the following auxiliary lemma:

▶ **Lemma 18.** *Let $n \in \mathbb{N}$ and $q \in (0,1)$. Then*

$$\sum_{k=n+1}^{\infty} k \cdot q^{k-1} \ =\ \frac{q^n}{(1-q)^2}\cdot(1+n-qn). \qquad (21)$$

Proof. By induction on n. For $n = 0$, the statement is equivalent to the well-known formula

$$\sum_{k=1}^{\infty} k\cdot(1-q)\cdot q^{k-1} \ =\ \frac{1}{1-q}$$

for the expected value of a geometrically distributed random variable. Next, assume that (21) holds for some $n \in \mathbb{N}$. Then

$$\sum_{k=n+2}^{\infty} k \cdot q^{k-1} = \sum_{k=n+1}^{\infty} k \cdot q^{k-1} - (n+1) \cdot q^n \overset{(21)}{=} \frac{q^n}{(1-q)^2} \cdot (1+n-qn) - (n+1) \cdot q^n$$
$$= \frac{q^n}{(1-q)^2} \cdot (1+n-qn-(n+1)\cdot(1-q)^2) = \frac{q^{n+1}}{(1-q)^2} \cdot (n+2-q(n+1)),$$

which is (21) for $n+1$. ◀

Now we are ready to prove Lemma 14:

Proof of Lemma 14. We compute

$$2\beta \cdot \sum_{i,j\in\mathbb{Z}_{\geq 0}:\max\{i,j\}>N} e^{-(i+j-1)\cdot\sigma\beta} \cdot \min\{B_i, B_j\}$$
$$\leq 4\beta \cdot \sum_{i=0}^{N} \sum_{j=N+1}^{\infty} e^{-(i+j-1)\cdot\sigma\beta} \cdot B_i + 2\beta \cdot \sum_{i=N+1}^{\infty} \sum_{j=N+1}^{\infty} e^{-(i+j-1)\cdot\sigma\beta} \cdot B_i$$
$$= 4\beta \cdot \sum_{i=0}^{N} e^{-(i-1)\cdot\sigma\beta} \cdot B_i \cdot \sum_{j=N+1}^{\infty} e^{-j\sigma\beta} + 2\beta \cdot \sum_{i=N+1}^{\infty} e^{-(i-1)\cdot\sigma\beta} \cdot B_i \cdot \sum_{j=N+1}^{\infty} e^{-j\sigma\beta}$$
$$= \delta_1 \cdot \sum_{i=0}^{N} e^{-(i-1)\cdot\sigma\beta} \cdot B_i + \frac{2\beta}{e^{N\sigma\beta}(e^{\sigma\beta}-1)} \cdot \sum_{i=N+1}^{\infty} e^{-(i-1)\cdot\sigma\beta} \cdot B_i$$

Bounding $B_i \leq i \cdot B_1$ for $i > N$ by using to the triangle inequality (Proposition 12), we obtain

$$\alpha \cdot \sum_{k=N+1}^{\infty} e^{-(k-1)\cdot\sigma\beta} \cdot B_k + 2\beta \cdot \sum_{i,j\in\mathbb{Z}_{\geq 0}:\max\{i,j\}>N} e^{-(i+j-1)\cdot\sigma\beta} \cdot \min\{B_i, B_j\}$$
$$\leq \delta_1 \cdot \sum_{i=0}^{N} e^{-(i-1)\cdot\sigma\beta} \cdot B_i + \left(\alpha + \frac{2\beta}{e^{N\sigma\beta}(e^{\sigma\beta}-1)}\right) \cdot \sum_{k=N+1}^{\infty} e^{-(k-1)\cdot\sigma\beta} \cdot B_k$$
$$\leq \delta_1 \cdot \sum_{i=0}^{N} e^{-(i-1)\cdot\sigma\beta} \cdot B_i + \left(\alpha + \frac{2\beta}{e^{N\sigma\beta}(e^{\sigma\beta}-1)}\right) \cdot \sum_{k=N+1}^{\infty} e^{-(k-1)\cdot\sigma\beta} \cdot k \cdot B_1$$
$$= \delta_1 \cdot \sum_{i=0}^{N} e^{-(i-1)\cdot\sigma\beta} \cdot B_i + \delta_2 \cdot B_1,$$

where we used Lemma 18 in the final equality with $n = N$ and $q = e^{-\sigma\beta}$. ◀

3 Discussion

We conjecture (but could not prove) that our lower bound examples (cf. [4]) are really worst-case examples, and that the values of our linear programs converge to these bounds.

Another question is whether the master route ratio is $\frac{1}{1-e^{-1/2}}$ even for low-activity instances. Currently we only know the upper bound of 3 from [25], but know no example with master route ratio larger than $\frac{1}{1-e^{-1/2}}$ (and this value is attained by our example only as the activity tends to infinity). The analogous question applies to the sampling algorithm: whether we need to consider the low-activity case separately is an open question.

Finally, we hope that our approach can also help for proving a better bound for related problems where similar random sampling techniques are used, or for showing that known bounds are best possible.

References

1 Mohamed Abdellahi Amar, Walid Khaznaji, and Monia Bellalouna. An exact resolution for the probabilistic traveling salesman problem under the a priori strategy. *Procedia Computer Science*, 108:1414–1423, 2017. doi:10.1016/j.procs.2017.05.068.

2 Dimitris Bertsimas. *Probabilistic combinatorial optimization problems*. PhD thesis, Massachusetts Institute of Technology, 1988. URL: https://dspace.mit.edu/handle/1721.1/14386.

3 Dimitris J. Bertsimas, Patrick Jaillet, and Amedeo R. Odoni. A priori optimization. *Operations Research*, 38(6):1019–1033, 1990. doi:10.1287/opre.38.6.1019.

4 Jannis Blauth, Meike Neuwohner, Luise Puhlmann, and Jens Vygen. Improved guarantees for the a priori TSP, 2023. doi:10.48550/arXiv.2309.10663.

5 Neill E. Bowler, Thomas M. A. Fink, and Robin C. Ball. Characterization of the probabilistic traveling salesman problem. *Phys. Rev. E*, 68:036703, 2003. doi:10.1103/PhysRevE.68.036703.

6 Martijn van Ee, Leo van Iersel, Teun Janssen, and René Sitters. A priori TSP in the scenario model. *Discrete Applied Mathematics*, 250:331–341, 2018. doi:10.1016/j.dam.2018.04.002.

7 Martijn van Ee and René Sitters. The a priori traveling repairman problem. *Algorithmica*, 80(10):2818–2833, 2018. doi:10.1007/s00453-017-0351-z.

8 Friedrich Eisenbrand, Fabrizio Grandoni, Thomas Rothvoß, and Guido Schäfer. Connected facility location via random facility sampling and core detouring. *Journal of Computer and System Sciences*, 76(8):709–726, 2010. doi:10.1016/j.jcss.2010.02.001.

9 Finn Fernstrøm and Teresa Anna Steiner. A constant approximation algorithm for the uniform a priori capacitated vehicle routing problem with unit demands. *Information Processing Letters*, 159-160:105960, 2020. doi:10.1016/j.ipl.2020.105960.

10 Arun Ganesh, Bruce M. Maggs, and Debmalya Panigrahi. Robust algorithms for TSP and Steiner Tree. *ACM Trans. Algorithms*, 19(2), 2023. doi:10.1145/3570957.

11 Naveen Garg, Anupam Gupta, Stefano Leonardi, and Piotr Sankowski. Stochastic analyses for online combinatorial optimization problems. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 942–951. SIAM, 2008. URL: https://dl.acm.org/doi/10.5555/1347082.1347185.

12 Michel Goemans and Jon Kleinberg. An improved approximation ratio for the minimum latency problem. *Mathematical Programming*, 82(1):111–124, 1998. doi:10.1007/BF01585867.

13 Igor Gorodezky, Robert D. Kleinberg, David B. Shmoys, and Gwen Spencer. Improved lower bounds for the universal and a priori TSP. In Maria Serna, Ronen Shaltiel, Klaus Jansen, and José Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 178–191. Springer, 2010. doi:10.1007/978-3-642-15369-3_14.

14 Anupam Gupta, Amit Kumar, Martin Pál, and Tim Roughgarden. Approximation via cost sharing: Simpler and better approximation algorithms for network design. *J. ACM*, 54(3), 2007. doi:10.1145/1236457.1236458.

15 Anupam Gupta, Amit Kumar, and Tim Roughgarden. Simpler and better approximation algorithms for network design. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, pages 365–372. ACM, 2003. doi:10.1145/780542.780597.

16 Anupam Gupta, Martin Pál, R. Ravi, and Amitabh Sinha. Boosted sampling: Approximation algorithms for stochastic optimization. In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, pages 417–426. ACM, 2004. doi:10.1145/1007352.1007419.

17 Patrick Jaillet. Probabilistic traveling salesman problems. PhD thesis, Massachusetts Institute of Technology, 1985. URL: https://dspace.mit.edu/handle/1721.1/15231.

18 Patrick Jaillet. A priori solution of a traveling salesman problem in which a random subset of the customers are visited. *Operations Research*, 36(6):929–936, 1988. doi:10.1287/opre.36.6.929.

19 Kamal Jain, Mohammad Mahdian, Evangelos Markakis, Amin Saberi, and Vijay V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing lp. *J. ACM*, 50(6):795–824, 2003. doi:10.1145/950620.950621.

20 Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. A (slightly) improved approximation algorithm for metric TSP. In *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing*, pages 32–45. ACM, 2021. doi:10.1145/3406325.3451009.

21 Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. A deterministic better-than-3/2 approximation algorithm for metric TSP. In Alberto Del Pia and Volker Kaibel, editors, *Integer Programming and Combinatorial Optimization*, pages 261–274. Springer, 2023. doi:10.1007/978-3-031-32726-1_19.

22 Fatemeh Navidi, Inge Li Gørtz, and Viswanath Nagarajan. Approximation algorithms for the a priori traveling repairman. *Operations Research Letters*, 48(5):599–606, 2020. doi:10.1016/j.orl.2020.07.009.

23 Christos H. Papadimitriou and Mihalis Yannakakis. The traveling salesman problem with distances one and two. *Mathematics of Operations Research*, 18(1):1–11, 1993. doi:10.1287/moor.18.1.1.

24 Frans Schalekamp and David B. Shmoys. Algorithms for the universal and a priori TSP. *Operations Research Letters*, 36(1):1–3, 2008. doi:10.1016/j.orl.2007.04.009.

25 David Shmoys and Kunal Talwar. A constant approximation algorithm for the a priori traveling salesman problem. In Andrea Lodi, Alessandro Panconesi, and Giovanni Rinaldi, editors, *Integer Programming and Combinatorial Optimization*, pages 331–343. Springer, 2008. doi:10.1007/978-3-540-68891-4_23.

26 Alejandro Toriello, William B. Haskell, and Michael Poremba. A dynamic traveling salesman problem with stochastic arc costs. *Operations Research*, 62(5):1107–1125, 2014. doi:10.1287/opre.2014.1301.

27 Laurence A. Wolsey. Heuristic analysis, linear programming and branch and bound. *Mathematical Programming Study*, 13:121–134, 1980. doi:10.1007/BFb0120913.

28 Anke van Zuylen. Deterministic sampling algorithms for network design. *Algorithmica*, 60(1):110–151, 2011. doi:10.1007/s00453-009-9344-x.