https://pubsonline.informs.org/journal/moor

MATHEMATICS OF OPERATIONS RESEARCH
Vol. 49, No. 2, May 2024, pp. 1169–1191
ISSN 0364-765X (print), ISSN 1526-5471 (online)

Probabilistic Bounds on the *k*-Traveling Salesman Problem and the Traveling Repairman Problem

Moïse Blanchard,[a,*] Alexandre Jacquillat,[a] Patrick Jaillet[a]

[a] Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

*Corresponding author

Contact: moiseb@mit.edu, https://orcid.org/0000-0003-0593-8666 (MB); alexjacq@mit.edu, https://orcid.org/0000-0002-2352-7839 (AJ); jaillet@mit.edu, https://orcid.org/0000-0002-8585-6566 (PJ)

Received: October 31, 2021
Revised: November 22, 2022
Accepted: May 16, 2023
Published Online in Articles in Advance: July 28, 2023

MSC2020 Subject Classifications: Primary: 90C27, 90B15; secondary: 60C05, 90B06

https://doi.org/10.1287/moor.2021.0286

Abstract. The k-traveling salesman problem (k-TSP) seeks a tour of minimal length that visits a subset of $k \leq n$ points. The traveling repairman problem (TRP) seeks a complete tour with minimal latency. This paper provides constant-factor probabilistic approximations of both problems. We first show that the optimal length of the k-TSP path grows at a rate of $\Theta\left(k/n^{k/(2(k-1))}\right)$. The proof provides a constant-factor approximation scheme, which solves a TSP in a high-concentration zone, leveraging large deviations of local concentrations. Then, we show that the optimal TRP latency grows at a rate of $\Theta(n\sqrt{n})$. This result extends the classic Beardwood–Halton–Hammersley theorem to the TRP. Again, the proof provides a constant-factor approximation scheme, which visits zones by decreasing order of probability density. We discuss practical implications of this result in the design of transportation and logistics systems. Finally, we propose dedicated notions of fairness—randomized population-based fairness for the k-TSP and geographic fairness for the TRP—and give algorithms to balance efficiency and fairness.

Funding: This work was partly supported by [Grant ONR N00014-18-1-2122] and the Singapore National Research Foundation through the Singapore–Massachusetts Institute of Technology Alliance for Research and Technology Centre for Future Urban Mobility.

Supplemental Material: The e-companion is available at https://doi.org/10.1287/moor.2021.0286.

Keywords: traveling salesman • stochastic model applications • suboptimal algorithms

1. Introduction

This paper studies the traveling repairman problem (TRP)—also known as the minimum latency problem (Afrati et al. [1], Bianco et al. [12], Minieka [30])—and the k-traveling salesman problem (k-TSP) in the Euclidean plane. These two problems are extensions of the well-studied traveling salesman problem (TSP). The TSP takes as inputs a set of n points as well as a distance matrix between all points and seeks the route of minimal length that visits all n points. Assuming constant speed, the TSP is equivalent to minimizing the arrival time at the end of the tour. Instead, the TRP seeks a tour that minimizes the sum of waiting times, known as the total latency. This problem arises in routing problems with requirements on customer wait times, for instance, to ensure sufficient level of service or to maximize operating profitability under random customer abandonment. The TRP is also applicable to disk head scheduling (Blum et al. [15]), flexible manufacturing systems (Simchi-Levi and Berman [38]), machine scheduling (Picard and Queyranne [35]), information search in computer networks (Ausiello et al. [4]) and other domains (Tsitsiklis [44]).

In contrast, the k-TSP seeks a path of minimal length that visits k out of n points, where $k \leq n$. In other words, the server chooses which points to serve. This problem has natural applications in routing and distribution systems, for example, for a logistics provider that can only serve a partial set of customers because of limitations on its delivery capacity. In addition, the k-TSP is used as subroutine for TRP approximation algorithms (Blum et al. [15], Goemans and Kleinberg [23]).

Our goal is to derive probabilistic bounds on the optimal k-TSP tour and the optimal TRP latency, which, in turn, lead to the design of efficient probabilistic approximation schemes. We consider a setting with a fixed number n of points in the Euclidean plane. The location of these points is unknown, following a known distribution; we denote by f the density of its absolutely continuous part. We seek constant-factor optimal approximations, that is, probabilistic solutions leading to an objective value that is asymptotically within a constant factor from the optimal solution. Specifically, we derive constant-factor estimates for the k-TSP and TRP solutions as a function of the number of points n and the density f. Moreover, through constructive proofs, we provide constant-factor approximation algorithms for both problems.

1.1. Related Work

The TSP is one of the canonical problems in operations research. The Beardwood–Halton–Hammersley (BHH) theorem, stated in Beardwood et al. [7] and improved in Steele [41, 42] gives a constant-factor $\Theta(\sqrt{n})$ approximation of the optimal TSP tour in the Euclidean space. The proof of these TSP estimates leads to the design of approximation algorithms that are stochastically robust in the a priori setting. A priori optimization (Bertsimas et al. [11]) provides an optimization framework when the same combinatorial problem is solved repeatedly over different instances. The goal is to compute a master solution ahead of time that minimizes an expected cost function given subsequent adjustments according to simple rules upon the realization of uncertainty.

This work has leveraged extensively the "locality property" of the TSP to design "divide and conquer" approximation algorithms. That is, under this approach, we define an a priori route that can then be slightly modified to respond to the instance realizations, keeping its approximation guarantees (Carlsson and Song [18]). Moreover, a near-optimal tour for the TSP objective remains near-optimal if we change the starting point of the tour. Even in the case of a unique starting depot, restricting the server to start serving from any point in the tour only induces an additional constant cost (which does not scale up with the number of points).

Despite its similarity with the TSP, the TRP lacks a locality property and is, therefore, much harder to solve. Local changes in the input points affect the waiting time of all the remaining ones, leading to nonlocal modifications in the optimal tour. Even in the one-dimensional case in which points lie on a line, the optimal TRP tour may cross itself several times, which is not the case in the TSP. Blum et al. [15] show that there exists a simple reduction from the TRP to the TSP, implying that the TRP is NP-hard in general for all metric spaces in which the TSP is known to be NP-hard. The TRP is even NP-hard on weighted trees, in which the TSP is easy (Sitters [39]).

Blum et al. [15] propose the first constant-factor approximation algorithm for the TRP in general metric spaces. Their approach involves a reduction to the k-minimum spanning tree (k-MST) problem, which seeks an optimal tree spanning k vertices in a weighted graph. This problem is also known to be NP-hard (Fischetti et al. [20]). Substantial work has been made to give approximation algorithms for this problem (Arya and Ramesh [3], Blum et al. [14], Garg [21], Ravi et al. [37]) with the current best bound being a 2-approximation algorithm (Garg [22]). More precisely, Blum et al. [15] show that a c-approximating algorithm for k-MST yields an $8c$-approximating algorithm for the TRP, thus providing a 16-approximation using the best known algorithm for the k-MST. Goemans and Kleinberg [23] improve the reduction in Blum et al. [15] from a factor of 8 to a factor of 3.59. Chaudhuri et al. [19] give the current best bound, a 3.59-approximation algorithm for the TRP in general metric spaces. In the case of weighted trees on the Euclidean plane, there exists a polynomial time $(1+\varepsilon)$-approximation algorithm (Sitters [40]).

The k-MST and k-TSP are also closely related. Hence, some papers on the k-MST give results for the k-TSP. Specifically, the algorithms given by Blum et al. [14], Garg [21, 22], and Arora and Karakostas [2] can be adapted to the k-TSP, which yields a 2-approximation algorithm for the k-TSP. These results are also leveraged to address other variants, such as prize-collector problems (Johnson et al. [26], Paul et al. [34]). More recently, Pandiri and Singh [33] give metaheuristics for the rooted k-TSP leveraging permutation-based and local-search heuristics.

Recent work focuses on the a priori TRP (Navidi et al. [31], van Ee and Sitters [45]). Following earlier work on the a priori TSP (Bertsimas et al. [11], Jaillet [25], Laporte et al. [27]), this problem seeks a master tour under demand uncertainty, in which each vertex is present with some probability. In this paper, we seek a priori solutions when the uncertainty lies in the position of the points as opposed to the number of such points.

Unlike the TSP, the k-TSP and the TRP encode a notion of priority between points. In the k-TSP, the decision maker can choose which points to serve; in the TRP, the decision maker can choose the sequence of customer visits. Such prioritization gives rise to important fairness issues. Namely, in the k-TSP, one can serve the points that lie in high-density zones, ignoring all other points altogether. Similarly, in the TRP, one can serve zones by decreasing order of density, thus prioritizing points in high-density zones over points in low-density zones. As a result, the approximation algorithms for both problems can lead to spatial discrimination across populations. This trade-off between efficiency and fairness arises in many resource allocation and scheduling problems (Bertsimas et al. [9, 10]), spanning communication networks (Bertsimas et al. [9], Luo et al. [28], Radunovic and Le Boudec [36]), air traffic management (Bertsimas and Gupta [8], Jacquillat and Vaze [24], Vossen et al. [46]), and finance (O'Cinneide et al. [32]).

1.2. Contributions and Outline

This paper makes three contributions:

– We derive a constant-factor probabilistic estimate of the optimal k-TSP tour for general distributions (Section 3). Specifically, we show that the optimal k-TSP length grows at a rate of $\Theta(k/n^{k/(2(k-1))})$. This result is obtained by leveraging large deviations in local point concentration to serve regions with high point concentration (especially for small k).

– We provide nonasymptotic constant-factor estimates of the optimal TRP for general distributions (Section 4). We show that total latency grows as $\Theta(n\sqrt{n})$ and characterize the dependence of the constant on the sampling distribution as the integral of a function of absolutely continuous part density, thus extending the BHH result from the TSP to the TRP. We discuss practical implications for the design of transportation and logistics systems in Section 2.3.

– We define fairness-enhanced versions of the k-TSP and TRP and analyze the price of fairness (Section 5). The approximation algorithms for the k-TSP and the TRP are highly "local." As a result, customers in high-density regions are more likely to receive a service (for the k-TSP) or to have a lower wait time (for the TRP). We define notions of fairness to circumvent this issue. For the TRP, we show that our approximation scheme satisfies max-min fairness and propose modifications toward proportional fairness. For the k-TSP, we show that geographical fairness across regions leads to significant efficiency loss. We, thus, propose population-based fairness given the distribution of populations across regions. We show that probabilistic population-based fairness still allows for flexibility and can lead to near-optimal k-TSP solutions.

Before proceeding, we first describe in Section 2 the modeling framework, outline our main results along with the proof techniques, and discuss their practical implications.

2. Setup, Overview of Results, and Practical Implications

2.1. Setup and Preliminaries

We consider a set of n points $V = \{X_1, \dots, X_n\}$ in the Euclidean space $\mathbb{R}^2$ equipped with the natural Euclidean distance. We focus on the two-dimensional case, but our results can easily be extended to the general case $\mathbb{R}^d$. We consider a probabilistic setting in which vertices $X_1, \dots, X_n$ are independent and identically distributed (i.i.d.), drawn from some distribution on a compact $\mathcal{K} \subset \mathbb{R}^2$. We denote by f the density of its absolutely continuous part.

Given the set of points V, we consider three optimization problems:

1. The TSP seeks a tour that starts in a vertex; visits all n vertices with some service order $x_1, \dots, x_n$; and returns to the starting point. The objective is to minimize the total length of the tour:

$$\sum_{i=1}^{n-1} |x_{i+1} - x_i| + |x_1 - x_n|. \tag{1}$$

2. The k-TSP seeks a path that visits an endogenous subset of $k \leq n$ vertices $x_1, \dots, x_k$. The objective is again to minimize the total length of the path:

$$\sum_{i=1}^{k-1} |x_{i+1} - x_i|. \tag{2}$$

3. The TRP, like the TSP, also seeks a complete tour of the n vertices. However, the TRP minimizes the total latency or the total wait times at the vertices. Formally, if $x_1, \dots, x_n$ defines a service order, the latency at point x_i is defined as $l_i = \sum_{j=1}^{i-1} |x_{j+1} - x_j|$. The TRP tour minimizes the sum of latencies:

$$\sum_{i=1}^{n} l_i = \sum_{i=1}^{n-1} (n-i)|x_{i+1} - x_i|. \tag{3}$$

In this paper, we provide constant-factor probabilistic bounds, that is, bounds on the expected optimal value of these problems that hold asymptotically within a universal constant factor, in which the expectation is taken over the randomness of the points $X_1, \dots, X_n$ (our bounds also hold with high probability). Similarly, we say that an algorithm is constant-factor optimal if it provides solutions with objective value within a constant factor of the optimal solution in expectation.

In this setting, the well-known BHH theorem shows that the optimal TSP length grows as $\Theta(\sqrt{n})$.

Theorem 1 (BHH Theorem, Beardwood et al. [7]). *Let $(X_i)_{i\geq 1}$ be a sequence of i.i.d. random points according to a distribution on a compact space $\mathcal{K} \subset \mathbb{R}^2$. With probability one, the length $l_{TSP}(X_1, \dots, X_n)$ of the optimal TSP on points $\{X_1, \dots, X_n\}$ satisfies*

$$\lim_{n\to\infty} \frac{l_{TSP}(X_1, \dots, X_n)}{\sqrt{n}} = \beta_{TSP} \iint_{\mathcal{K}} \sqrt{f(x)}dx,$$

where $0.6250 \leq \beta_{TSP} \leq 0.9204$ is a universal constant and f denotes the density of the absolutely continuous part of the distribution.

Lemma 1 provides a simplified version of Theorem 1 that is useful in our analysis. The proof of this result constructs a simple "master" space-filling curve that is at most $1/(2\sqrt{n})$ away from any point in the unit square and has length $\sqrt{n} + \mathcal{O}(1)$. We can adapt this simple curve to serve any vertex by adding a back-and-forth detour from the closest point on the curve. Similarly, we can adapt the curve to serve n points. The length of the resulting tour is $2\sqrt{n} + \mathcal{O}(1)$.

Lemma 1 (Beardwood et al. [7]). *Let $n \geq 2$ and $(X_i)_{1\leq i\leq n}$ points in the unit square $[0,1]^2$. Denote by $l_{TSP}(X_1, \ldots, X_n)$ the length of the TSP tour visiting these points. Then,*

$$l_{TSP}(X_1, \ldots, X_n) \leq 2\sqrt{n} + C,$$

for some universal constant $C > 0$.

In our algorithms for the k-TSP and the TRP, we use this result as a subroutine to design an a priori curve that can serve n points with a worst case length of $2\sqrt{n} + C$. Asymptotically, this a priori procedure yields solutions that are at most $2/\beta_{TSP}$ away from the optimal TSP tour.

2.2. Main Results

The main results of the paper provide constant-factor approximations of the k-TSP and TRP solutions. First, we show in Section 3 that the optimal k-TSP tour grows at a rate of $\Theta\left(k/n^{k/(2(k-1))}\right)$ (Theorem 2). This rate can be interpreted as a positive result by contrasting it with (i) a naive bound of $\sqrt{k}$, which applies a TSP tour on a random subset of k points, and (ii) a bound of $k/\sqrt{n}$, which selects the best subpath of k consecutive vertices in the full TSP tour. The rate of $\Theta\left(k/n^{k/(2(k-1))}\right)$ underscores a benefit of $\sqrt{k/n}$ that comes from merely optimizing which vertices to visit and an additional benefit of $n^{-1/(2(k-1))}$ that comes from reoptimizing the tour, leveraging large deviations in local point concentration.

The proof of the k-TSP proceeds by showing that the rate $k/n^{k/(2(k-1))}$ is nonasymptotically tight up to a constant with uniform densities. We extend the analysis to the case of general measurable (not necessarily continuous) densities. In particular, the proof for the upper bound is constructive and provides a constant-factor approximation algorithm when $1 \ll k \ll n$ by selecting the region with highest point concentration and performing the (uniform) k-TSP in this region.

Second, we show in Section 4 that the optimal TRP latency grows at a rate of $\Theta(n\sqrt{n})$ (Theorem 3). In contrast to the previous one, this is a rather negative result. Indeed, the TSP tour gives a $\Theta(\sqrt{n})$ estimate of the latency in the last vertex. Accordingly, if all customers had to wait as long as the last customer, we would end up with a total latency of the order of $n\sqrt{n}$. As this result shows, even by reoptimizing the tour, the TRP still leads to optimal latency on the order of $n\sqrt{n}$.

The proof of the TRP upper bound is also constructive and gives a simple constant-factor approximation scheme. This scheme constructs a master a priori tour, depending solely on the absolutely continuous part density and then adapts it to any realization of sampled points. Specifically, the algorithm partitions the region into zones of constant density, visits zones by decreasing order of local density, and performs a tour on each zone following space-filling techniques for the TSP.

From a practical standpoint, the TRP result is structurally different from the TSP result. Specifically, the optimal TSP tour is concave in the number of vertices, indicating economies of scale. In contrast, the optimal TRP latency is convex in the number of vertices, indicating diseconomies of scale. This distinction has implications for the design of transportation and logistics systems.

2.3. Implications for Transportation and Logistics Operations

TSP approximation results provide insights into the operations of transportation and logistics systems, which can be used to support upstream planning decisions. Sample applications include location analysis (Carlsson and Jones [17]), area partitioning for vehicle routing (Carlsson [16]), and same-day delivery systems (Banerjee et al. [5, 6], Stroh et al. [43]). In these problems, continuous approximations estimate routing costs into upstream optimization models rather than, for instance, capturing discrete routing dynamics at significant computational costs.

Specifically, TSP approximation results take the perspective of a logistics provider. However, several systems strive to also minimize customer wait times. For instance, in food delivery, a company needs to serve customers as early as possible as opposed to meeting an overall deadline. As another example, school bus (or company bus) routing aims to minimize the travel times of the students (or employees) as opposed to the vehicle's trip time. The TRP provides the natural framework to estimate customer level of service. As such, the results of this paper can be used to guide the design of such transportation and logistics systems focused on wait times.

This distinction between the TSP length and TRP latency has practical consequences because of the concavity of the $\sqrt{n}$ function versus the convexity of the $n\sqrt{n}$ function. As a result, economies of scale in the TSP favor service concentration (few vehicles, each serving many customers), whereas diseconomies of scale in the TRP favor service dispersion (more vehicles each serving a smaller number of customers). We illustrate this tension in two simple examples.

2.3.1. Fleet Size Optimization. We seek the number of vehicles m to serve a batch of N orders. Each vehicle incurs a fixed cost c and carries N/m orders. Assume first that the system minimizes vehicles' fixed and travel costs. Based on the BHH approximation, we can write this objective as minimizing $c\cdot m + d\cdot m\cdot\sqrt{N/m} = cm + d\sqrt{Nm}$ for some scaling constant d. The optimal strategy is $m=1$ even with $c=0$; that is, a single vehicle serves all customers. However, if we replace the vehicle travel time component with a customer wait time component, the objective becomes minimizing $c\cdot m + \tilde{d}\cdot N/m\sqrt{N/m}$ for some scaling constant $\tilde{d}$. The optimum is now attained for $m^* = (3\tilde{d}/(2c))^{2/5}N^{3/5}$. Now, the operator leverages a multivehicle fleet, which increases with customer demand. This example underscores two opposite strategies, spanning pure consolidation in the TSP case (serving the entire batch with a single vehicle) versus dispersion in the TRP case (serving customer demand with multiple vehicles to balance vehicle costs and customer wait times).

2.3.2. Vehicle Dispatch in Same-Day-Delivery (SDD) Systems. Based on Stroh et al. [43], we consider an SDD provider that operates a fleet of m vehicles, each of which can only be dispatched once. Customers arrive at a constant rate λ until an order cutoff N is met at time $T_{cutoff} = N/\lambda$. The operator optimizes dispatch decisions, characterized by a dispatch time t_i and a number of carried orders n_i for each vehicle $i=1,\ldots,m$. Following the BHH approximation, the delivery time of vehicle i can be written as $a\cdot\sqrt{n_i}$ for some scaling constant a. The SDD constraint asks that vehicles should complete their deliveries by an end-of-day deadline T, that is, $t_i + a\sqrt{n_i} \le T$ for all $i=1,\ldots,m$. Stroh et al. [43] minimize the total dispatch time $\sum_{i=1}^m a\sqrt{n_i}$ under the aforementioned SDD, demand (all orders need to be served), and consistency (orders can only be carried after they become available) constraints. Whenever feasible, the optimal strategy is to dispatch the first vehicle when it can fulfill all revealed orders and return exactly at time T, the second vehicle when it can fulfill all subsequent orders and return exactly at time T, etc. (top of Figure 1). This strategy is feasible (hence, optimal) whenever the fleet m is sufficiently large to cover all the demand, which can be checked by solving recursively the equations $t_i + a\sqrt{\lambda(t_i - t_{i-1})} = T$ for $t_{i-1} \le t_i \le T$ with $t_0 = 0$ and checking whether $t_m \ge T_{cutoff}$.

Now, assume that the operator minimizes customer wait times. Based on our TRP approximation result, this scales as $w\cdot n\sqrt{n}$ for some scaling constant w. Note that the cost function can be augmented by replacing $wn\sqrt{n}$ with $b\cdot n^2 + w\cdot n\sqrt{n}$, where $b\cdot n^2$ captures the batching time prior to the dispatch and $w\cdot n\sqrt{n}$ captures the wait time after the dispatch. Either way, the cost function is now convex in n. Whenever feasible, the optimal strategy is, therefore, to dispatch vehicles at regular times $iN/(m\lambda)$ (bottom of Figure 1). This strategy is feasible (hence, optimal) whenever the last vehicle m can complete its orders by the end of the day, that is, whenever $N/\lambda + a\sqrt{N/m} \le T$.

Again, this structure underscores two opposite strategies. In the SDD system (based on a TSP objective), the dispatching policy leverages consolidation by bundling orders together as much as possible. In contrast, in the food

Figure 1. (Color online) Consolidation-driven dispatch based on order deadlines from the TSP approximation (top) versus dispersion-driven dispatch based on customer wait times from the TRP approximation (bottom) for $m = 4$ vehicles.

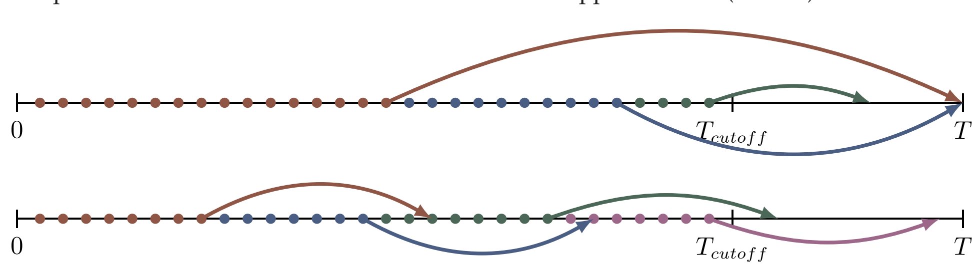

delivery, school bus, and employee bus systems, the dispatching policy leverages dispersion by distributing orders as evenly as possible. Although stylized, these two examples underscore that minimizing wait times may significantly alter design decisions in routing systems as compared with focusing on vehicle travel times.

3. The *k*-Traveling Salesman Problem

We provide probabilistic estimates on the length of the *k*-TSP tour. Before proceeding, let us expand on the two aforementioned naive bounds:

– Upper bound of $\mathcal{O}(\sqrt{k})$: By choosing the k points to visit uniformly at random among the n available points, the BHH theorem ensures that the length of the optimal path visiting these k points has length $\sim \beta_{TSP}\sqrt{k}\iint_{\mathcal{K}}\sqrt{f}$ as $k \to \infty$. However, this analysis does not leverage the flexibility regarding which points to serve.

– Upper bound of $\mathcal{O}(k/\sqrt{n})$: Consider the optimal TSP tour visiting all n points of length $l_{TSP}(n)$. Selecting k consecutive points on this tour at random—we randomly select the starting point—yields a path of length $(k-1)/n \cdot l_{TSP}(n)$ in expectation. In particular, the best choice of k consecutive points on the TSP tour yields an upper bound for the k-TSP of $(k-1)/n \cdot l_{TSP}(n) = \mathcal{O}(k/\sqrt{n})$. This observation underscores the benefits of choosing which points to serve. As we see, such flexibility can be very significant, especially for small values of k. Yet this analysis still relies on the optimal TSP tour, therefore eliminating an extra degree of freedom in the k-TSP.

We show that this rate $\mathcal{O}(k/\sqrt{n})$ is essentially tight for large k but can be tightened for small k. For instance, for $k=2$, the minimum distance between n uniformly sampled points in the unit square is $\Theta(1/n)$ instead of $\mathcal{O}(1/\sqrt{n})$. Our results in this section interpolate the $\Theta(1/n)$ estimate for $k=2$ and the $\Theta(\sqrt{n})$ estimate for $k=n$. We now present the main result of this section giving the exact rate of the expected k-TSP length. Note that this result does not only provide an asymptotic rate, but holds yields an estimate of the k-TSP length for any choice of $2 \le k \le n$.

Theorem 2. *Assume n vertices are drawn independently, uniformly on a compact space $\mathcal{K} \subset \mathbb{R}^2$ with area $\mathcal{A}_{\mathcal{K}}$. Denote by $l_{TSP}(k,n)$ the length of the k-TSP on these n vertices. Then, for all $n \ge 2$ and $2 \le k \le n$, for some universal constants $0 < c < C$,*

$$c\frac{k-1}{n^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}}\sqrt{\mathcal{A}_{\mathcal{K}}} \le \mathbb{E}[l_{TSP}(k,n)] \le C\frac{k-1}{n^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}}\sqrt{\mathcal{A}_{\mathcal{K}}}.$$

Theorem 2 exhibits an additional factor $\Theta(n^{-1/(2(k-1))})$ compared with the previous bound $\mathcal{O}(k/\sqrt{n})$. This additional factor corresponds to large deviations of local point densities. Consider any subsquare of area $\mathcal{O}(k/n)$ and perform the TSP on this subsquare. We expect $\mathcal{O}(k)$ points in this subsquare, yielding a path of length $\mathcal{O}(\sqrt{k} \cdot \sqrt{k/n}) = \mathcal{O}(k/\sqrt{n})$. In the k-TSP however, we can choose to serve zones with abnormally high point concentration, deviating from the expected density. In the following two sections, we prove Theorem 2 and show that the resulting discount on the length of the optimal path visiting k points is the additional factor $\Theta(n^{-1/(2(k-1))})$. A comparison of the obtained convergence rate together with a comparison to the simple bound $\mathcal{O}(k/\sqrt{n})$ is represented in Figure 2.

Figure 2. (Color online) Convergence rate of the length of the k-TSP $\Theta\left((k-1)/n^{k/(2(k-1))}\right)$ (in solid lines) compared with the rate of convergence of the simple heuristic $\Theta((k-1)/\sqrt{n})$ (in dashed lines) as a function of n and k.

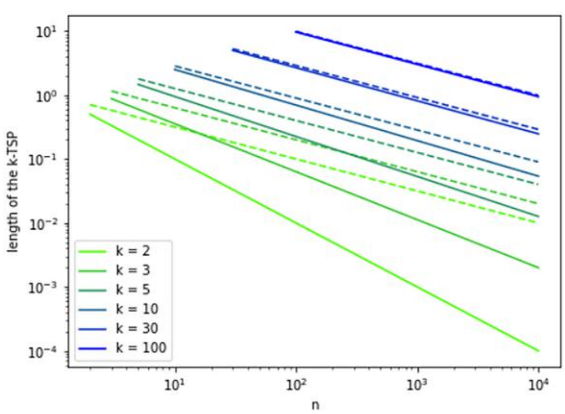

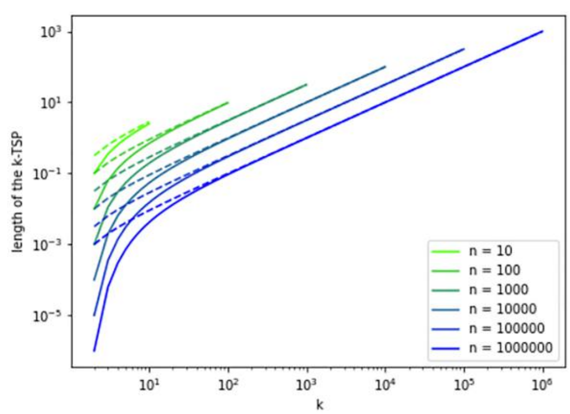

3.1. Lower Bounds on the *k*-TSP

We first need the following lemma.

Lemma 2. *Assume all n vertices are drawn independently, uniformly on a compact $\mathcal{K}\subset\mathbb{R}^2$ with area $\mathcal{A}_\mathcal{K}$. Denote by $l_{TSP}(k,n)$ the length of the k-TSP on these n vertices. Then, for any $\alpha>0$,*

$$\mathbb{P}[l_{TSP}(k,n)\le\alpha]\le n^k\left(\frac{2\pi\alpha^2}{\mathcal{A}_\mathcal{K}}\right)^{k-1}\frac{1}{(2k-2)!}.$$

Proof. By symmetry on the vertices and because $n!/((n-k)!)\le n^k$,

$$\begin{aligned}\mathbb{P}[l_{TSP}(k,n)\le\alpha]&=\mathbb{E}[\mathbf{1}_{l_{TSP}(k,n)\le\alpha}]\\&\le\mathbb{E}\left[\sum_{1\le i_1,\dots,i_k\le n\text{ distinct}}\mathbf{1}(|v_{i_2}-v_{i_1}|+\cdots+|v_{i_k}-v_{i_{k-1}}|\le\alpha)\right]\\&\le n^k\mathbb{E}[\mathbf{1}_{|v_2-v_1|+\cdots+|v_k-v_{k-1}|\le\alpha}].\end{aligned}$$

We next estimate the last term. Given the position of v_1, the probability of having $l_1\le|v_2-v_1|\le l_1+dl_1$ is at most $(2\pi l_1/\mathcal{A}_\mathcal{K})dl_1$. Similarly, conditionally on $v_1,\dots v_{k-1}$, the probability of having $l_{k-1}\le|v_k-v_{k-1}|\le l_{k-1}+dl_{k-1}$ is at most $(2\pi l_{k-1}/\mathcal{A}_\mathcal{K})dl_{k-1}$ (see Figure 3 for an illustration for $k=4$). Therefore,

$$\mathbb{E}[\mathbf{1}_{|v_2-v_1|+\cdots+|v_k-v_{k-1}|\le\alpha}]\le\int_{l_1,\dots,l_{k-1}\ge0}\mathbf{1}_{l_1+\cdots+l_{k-1}\le\alpha}\left(\frac{2\pi l_1}{\mathcal{A}_\mathcal{K}}\right)\cdots\left(\frac{2\pi l_{k-1}}{\mathcal{A}_\mathcal{K}}\right)dl_1\cdots dl_{k-1}=\left(\frac{2\pi\alpha^2}{\mathcal{A}_\mathcal{K}}\right)^{k-1}\mathcal{P}_{k-1},$$

where $\mathcal{P}_{k-1}:=\int_{l_1,\dots,l_{k-1}\ge0}\mathbf{1}_{l_1+\cdots+l_{k-1}\le1}\cdot l_1\cdots l_{k-1}\cdot dl_1\cdots dl_{k-1}$. Now, for any $k\ge2$,

$$\begin{aligned}\mathcal{P}_k&=\int_{l_k=0}^1 l_k\left(\int_{l_1,\dots,l_{k-1}\ge0}\mathbf{1}_{l_1+\cdots+l_{k-1}\le1-l_k}\cdot l_1\cdots l_{k-1}\cdot dl_1\cdots dl_{k-1}\right)dl_k\\&=\int_0^1 l_k\cdot(1-l_k)^{2(k-1)}\mathcal{P}_{k-1}\cdot dl_k=\mathcal{P}_{k-1}\cdot\frac{1}{(2k-1)(2k)}.\end{aligned}$$

Because $\mathcal{P}_1=\frac{1}{2}$, by induction $\mathcal{P}_k=1/((2k)!)$. Putting everything together yields the desired result. □

We are now ready to prove a lower bound on the *k*-TSP.

Proof of the Lower Bound in Theorem 2. Applying Lemma 2, we obtain

$$\begin{aligned}\mathbb{P}\left[l_{TSP}(k,n)\le\varepsilon\sqrt{\frac{2}{e^2\pi}\frac{k-1}{n^{\frac{1}{2}(1+\frac{1}{k-1})}}}\sqrt{\mathcal{A}_\mathcal{K}}\right]&\le n^k\left(\frac{4\varepsilon^2(k-1)^2}{e^2\cdot n^{(1+\frac{1}{k-1})}}\right)^{k-1}\frac{1}{(2k-2)!}\\&\le\left(\frac{4\varepsilon^2(k-1)^2}{e^2}\right)^{k-1}\frac{1}{2\sqrt{\pi(k-1)}}\left(\frac{e}{2(k-1)}\right)^{2(k-1)}\\&=\frac{\varepsilon^{2k-2}}{2\sqrt{\pi(k-1)}},\end{aligned}$$

Figure 3. (Color online) Illustration of the proof of Lemma 2: $\mathbb{P}(l_i\le|v_{i+1}-v_i|\le l_i+dl_i)\le(2\pi l_i/\mathcal{A}_\mathcal{K})dl_i$.

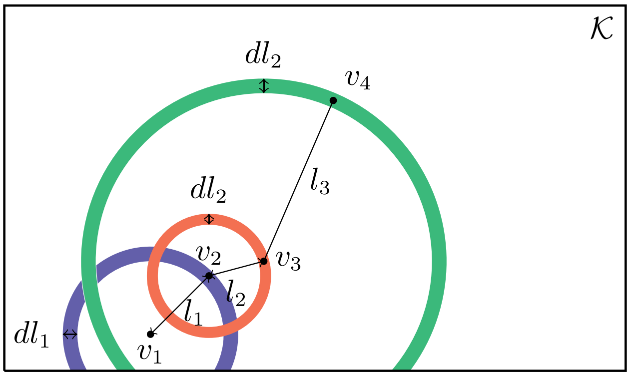

where we use Stirling's approximation $\sqrt{2\pi}n^{n+1/2}e^{-n} \le n! \le en^{n+1/2}e^{-n}$. Then,

$$\begin{aligned}
\mathbb{E}[l_{TSP}(k,n)] &= \sqrt{\frac{2}{e^2\pi}}\frac{k-1}{n^{\frac{1}{2}(1+\frac{1}{k-1})}}\sqrt{\mathcal{A}_\mathcal{K}}\int_0^\infty \mathbb{P}\left[l_{TSP}(k,n) \ge \varepsilon\sqrt{\frac{2}{e^2\pi}}\frac{k-1}{n^{\frac{1}{2}(1+\frac{1}{k-1})}}\sqrt{\mathcal{A}_\mathcal{K}}\right]d\varepsilon \\
&\ge \sqrt{\frac{2}{e^2\pi}}\frac{k-1}{n^{\frac{1}{2}(1+\frac{1}{k-1})}}\sqrt{\mathcal{A}_\mathcal{K}}\int_0^1 \left(1 - \frac{\varepsilon^{2k-2}}{2\sqrt{\pi(k-1)}}\right)d\varepsilon \\
&\ge \sqrt{\frac{2}{e^2\pi}}\frac{k-1}{n^{\frac{1}{2}(1+\frac{1}{k-1})}}\sqrt{\mathcal{A}_\mathcal{K}}\left(1-\frac{1}{6\sqrt{\pi}}\right),
\end{aligned}$$

where in the last inequality, we use $\int_0^1 \varepsilon^{2k-2}/\sqrt{k-1} \le \int_0^1 \varepsilon^2 = 1/3$. The result follows. □

This lower bound improves over the simple rate $\mathcal{O}(k/\sqrt{n})$ obtained by using the TSP tour only. In particular, when k is small, we can improve the exponent of the denominator; for example, for $k=1$, we obtain the rate $\Omega(1/n)$, and for $k=2$, we get a rate $\Omega(1/n^{3/4})$. For $k=\Omega(\log n)$, the term $1/(k-1)$ in the exponent of the denominator can be omitted. Thus, the provided lower bound becomes $\Omega(k/\sqrt{n})$, matching the simple upper bound with high probability as shown in the following result.

Corollary 1. *Assume all n vertices are drawn independently, uniformly on a compact space $\mathcal{K} \subset \mathbb{R}^2$ with area $\mathcal{A}_\mathcal{K}$. Denote by $l_{TSP}(k,n)$ the length of the k-TSP on these n vertices. Then, there exists a universal constant $M > 0$ such that, for $M \log n \le k_n \le n$,*

$$\mathbb{P}\left[l_{TSP}(k_n,n) \le \frac{k_n}{e\sqrt{\pi n}}\sqrt{\mathcal{A}_\mathcal{K}}\right] = o(e^{-k_n}).$$

Proof. We use Lemma 2 with $\alpha = k_n\sqrt{\mathcal{A}_\mathcal{K}}/(e\sqrt{\pi n})$ and the lower bound $\sqrt{2\pi}n^{n+1/2}e^{-n} \le n!$ to get

$$\mathbb{P}\left[l_{TSP}(k_n,n) \le \frac{k_n}{e\sqrt{\pi n}}\sqrt{\mathcal{A}_\mathcal{K}}\right] \le n\left(\frac{2k_n^2}{e^2}\right)^{k_n-1}\frac{4k_n^2}{(2k_n)!} \le \frac{4e^2 n}{\sqrt{\pi k_n}2^{k_n}} \le \frac{4e^2}{\sqrt{\pi}}\cdot e^{\log n - k_n \log 2}.$$

Therefore, for $M > 2/\log 2$ and for all $k_n \ge M\log n$, the right-hand side term is $o(e^{-k_n})$. □

3.2. Upper Bound on the *k*-TSP

In this section, we show that the lower bound shown in Section 3.1 is tight up to a constant factor.

Proof of the Upper Bound of Theorem 2. We first suppose $k \le n^{1/3}$ and treat the case $k \ge n^{1/3}$ separately. Fix $\alpha > 0$. We start by covering the compact $\mathcal{K}$ into P_α disjoint subsquares of equal size $(1/m_\alpha) \times (1/m_\alpha)$, where

$$m_\alpha := \left\lfloor \frac{1}{\alpha}\sqrt{n^{1+\frac{1}{k-1}}/(\mathcal{A}_\mathcal{K}(k-1))}\right\rfloor.$$

Because $\mathcal{K}$ is measurable and has area $\mathcal{A}_\mathcal{K}$, we know that $P_\alpha \sim \mathcal{A}_\mathcal{K} m_\alpha^2$ as $\alpha \to \infty$. We first show that, with high probability, there exists at least one of these subsquares that contains at least k vertices, and we upper bound $l_{TSP}(k,n)$ by the length of the TSP tour in that subsquare (see Figure 4). Define X_i^α as the number of vertices in subsquare i for $1 \le i \le P_\alpha$. Then, $(X_1^\alpha, \ldots, X_{P_\alpha}^\alpha)$ follows a multinomial distribution with n trials and uniform probabilities $1/P_\alpha$. Denote by $A_i^\alpha = \{X_i^\alpha \ge k\}$ the event that subsquare i contains at least k vertices. For any $1 \le i \le P_\alpha$, using the fact that $P_\alpha = o(1/n)$,

$$\begin{aligned}
\mathbb{P}(A_i^\alpha) = \mathbb{P}(A_1^\alpha) &\ge \binom{n}{k}\frac{1}{{P_\alpha}^k}\left(1-\frac{1}{P_\alpha}\right)^{n-k} \\
&\ge \frac{1}{k!}\cdot\frac{n^k}{{P_\alpha}^k}\cdot(1+o(1)) \\
&\ge \frac{(1+o(1))^k(k-1)^k}{k!}\frac{\alpha^{2k}}{n^{1+\frac{1}{k-1}}}\cdot(1+o(1)) \\
&\ge c\cdot\frac{\alpha^{2k-2}}{P_\alpha},
\end{aligned}$$

Figure 4. Illustration of the proof of the upper bound of Theorem 2 for $n=20$ and $k=7$. The procedure partitions the square into subsquares and then performs the TSP on k points in a subsquare containing at least k points.

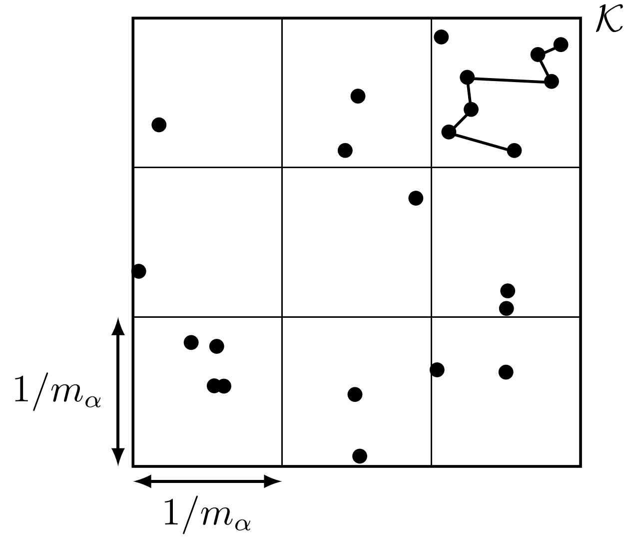

for some constant $c>0$. Then, the Bonferroni–Mallows bound for multinomials (Mallows [29]) implies

$$\begin{aligned}\mathbb{P}\left[\bigcup_{i=1}^{P_\alpha} A_i^\alpha\right] &= 1-\mathbb{P}[X_1^\alpha \le k-1,\ldots,X_{P_\alpha}^\alpha \le k-1]\\ &\ge 1-\prod_{i=1}^{P_\alpha}\mathbb{P}(X_i^\alpha \le k-1)\\ &\ge 1-e^{-\sum_{i=1}^{P_\alpha}\mathbb{P}(A_i^\alpha)} \ge 1-e^{-c\cdot\alpha^{2k-2}}.\end{aligned}$$

Now, assume that the event $\cup_{i=1}^{P_\alpha} A_i^\alpha$ is met. Let $1\le i\le P_\alpha$ be the index of a subsquare that contains at least k vertices. Then, according to Lemma 1, the length of the TSP on any k vertices in this subsquare of size $(1/m_\alpha)\times(1/m_\alpha)$ is at most $(2\sqrt{k}+C)/m_\alpha \le \tilde{C}\alpha(k-1)\sqrt{\mathcal{A}_\mathcal{K}}/n^{k/(2(k-1))}$ for some universal constant $\tilde{C}>0$. Therefore, using the previous equation, we get

$$\mathbb{P}\left[l_{TSP}(k,n) > \tilde{C}\alpha\frac{k-1}{n^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}}\sqrt{\mathcal{A}_\mathcal{K}}\right] \le 1-\mathbb{P}\left[\bigcup_{i=1}^{P_\alpha} A_i^\alpha\right] \le e^{-c\cdot\alpha^{2k-2}}.$$

Finally, we apply this inequality to obtain

$$\begin{aligned}\mathbb{E}[l_{TSP}(k,n)] &\le \tilde{C}\frac{k-1}{n^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}}\sqrt{\mathcal{A}_\mathcal{K}} + \int_{\tilde{C}(k-1)/n^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}\sqrt{\mathcal{A}_\mathcal{K}}}^{\infty}\mathbb{P}[l_{TSP}(k,n)>x]dx\\ &\le \tilde{C}\frac{k-1}{n^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}}\sqrt{\mathcal{A}_\mathcal{K}}\left(1+\int_1^\infty e^{-c\cdot\alpha^{2k-2}}d\alpha\right) \le \hat{C}\frac{k-1}{n^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}}\sqrt{\mathcal{A}_\mathcal{K}},\end{aligned}$$

for some universal constant $\hat{C}$. This ends the proof for $k\le n^{1/3}$. Now, consider the case $k\ge n^{1/3}$. In this case, $n^{k/(2(k-1))}\sim\sqrt{n}$, hence, the result can be derived from TSP bounds: let l_k^* be the minimum length of a subpath of the optimal TSP tour with k consecutive vertices. Because the average length of a path visiting k consecutive vertices is exactly $(k-1)/n\cdot l_{TSP}$, Theorem 1 yields directly $\mathbb{E}[l_{TSP}(k,n)]\le\mathbb{E}[l_k^*]\lesssim(k-1)/n\cdot\beta_{TSP}\sqrt{n\mathcal{A}_\mathcal{K}}$. □

The proof of the upper bound is constructive and, therefore, gives a simple algorithm reaching this bound: first, partition the unit square into P_α equal subsquares, select a subsquare with at least k points, then perform the TSP on any k points in this subsquare (see Figure 4). There exists such a subsquare with very high probability. To obtain a constant-factor approximation, we only need a constant-factor approximation of the TSP in the subsquare. For instance, we can use the simple procedure from Lemma 1 to obtain a path of length at most $(2\sqrt{k}+\mathcal{O}(1))/\sqrt{P_\alpha}$. This procedure may fail to produce a path if no subsquare contains k points, but one can repeat the procedure successively for $\alpha = 1, 2, 3\ldots$ until we find a subsquare with at least k points. By Theorem 2, this algorithm is a constant-factor approximation to the k-TSP in expectation.

3.3. Generalization to Nonuniform Distributions

Theorem 2 may be generalized to the case in which point positions are drawn independently according to some distribution with a density f. For simplicity, we suppose that the density is continuous, but the result can be extended to more general densities via smoothing techniques (e.g., Lebesgue derivatives); this is detailed in a companion report (Blanchard et al. [13]). Because the density is continuous, we can focus on the region of maximum density $\|f\|_\infty$ and relate the k-TSP on n points sampled with f to the k-TSP on $\|f\|_\infty n$ points sampled uniformly. Hence, we expect the guarantees of Theorem 2 to hold, replacing n with $\|f\|_\infty n$.

Proposition 1. *Assume n vertices are drawn independently on a compact space $\mathcal{K}$, according to a continuous density f. Denote by $l_{TSP}(k, n)$ the length of the k-TSP on these n vertices, where $2 \le k \le n$. There exists a universal constant $c > 0$ such that*

$$\liminf_{n\to\infty} \mathbb{E}[l_{TSP}(k,n)] \frac{(\|f\|_\infty n)^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}}{k-1} \mathcal{A}_{\mathcal{K}}^{\frac{1}{2(k-1)}} \ge c,$$

and, further, if $k/n \to 0$ and $k \to \infty$ as $n \to \infty$, there exists a universal constant $C > 0$ such that

$$\limsup_{n\to\infty} \mathbb{E}[l_{TSP}(k,n)] \frac{(\|f\|_\infty n)^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}}{k-1} \le C.$$

Proof. For the lower bound, we use a standard sample-and-reject argument to upper sample the n points according to f from the uniform density on $\mathcal{K}$ as follows. Consider a sequence (X_i) of i.i.d. uniformly drawn points. A point $X_i = x_i$ is rejected independently of the other points with probability $1 - f(x_i)/\|\tilde{f}\|_\infty$. The sequence (Y_i) is i.i.d. distributed according to f. Using the Hoeffding inequality, we show that, with probability $1 - e^{-n^2/2}$, from $N := \lceil 2\|\tilde{f}\|_\infty \mathcal{A}_{\mathcal{K}} n\rceil$ uniform draws $(X_i)_{i\le N}$, at least n points are drawn according to density f with the rejection process. On this event, we lower bound the k-TSP length on n points drawn according to f with the k-TSP length on the N vertices $(X_i)_{i\le N}$. Therefore, using Theorem 2, for some constant $\tilde{c} > 0$,

$$\mathbb{E}[l_{TSP(f)}(k,n)] \ge (1 - e^{-n^2/2})\mathbb{E}[l_{TSP(\mathcal{U})}(k,N)] \gtrsim \frac{\tilde{c}}{2} \frac{k-1}{(\|\tilde{f}\|_\infty n)^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}} \mathcal{A}_{\mathcal{K}}^{-\frac{1}{2(k-1)}}.$$

Therefore, we obtain the desired lower bound. For the upper bound, because f is continuous, there exists a nonempty square U such that the density is at least $\|f\|_\infty/2$ on U. By the Hoeffding inequality, with probability at least $1 - e^{-\varepsilon^2(\|f\|_\infty^2/2)\mathcal{A}_U^2 n}$, at least $n_U = \|f\|_\infty \mathcal{A}_U(1-\varepsilon)n/2$ points fell in U. Denote by E_0 this event on which these n_U vertices are drawn uniformly on U. Then, using Theorem 2,

$$\mathbb{E}[l_{TSP(f)}(k,n)] \le n\mathrm{diam}(\mathcal{A}_{\mathcal{K}})\mathbb{P}[E_0^c] + \mathbb{E}[l_{TSP(U)}(k,n_U)] \le (1 + o_n(1)) \cdot C \frac{k-1}{n_U^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}} \sqrt{\mathcal{A}_U}.$$

Because $\mathcal{A}_U^{1/(k-1)} = \mathcal{O}(1)$, the desired upper bound follows. □

The intuition of this generalization is fairly simple: instead of solving the k-TSP on the whole compact space $\mathcal{K}$, we can focus on zones in which the density is maximal. The hypothesis $k_n = o(n)$ ensures that this restriction is feasible (otherwise, there would not be k_n points locally). When $k = o(n)$ and $k \to \infty$, the proposed local strategy—performing the k-TSP on the highest density zone—is constant-factor optimal in expectation. As suggested by Proposition 1, this is not exactly the case when $k = \mathcal{O}(1)$, for which restricting to a fixed high-density zone affects the local concentration property of the large-deviations analysis.

4. The Traveling Repairman Problem

We now turn to the TRP, which seeks a tour minimizing total latency (Equation (3)). For simplicity, assume that we can choose any point as the starting point. Indeed, we show that the TRP objective is $\Theta(n\sqrt{n})$, whereas an initial edge from a fixed depot to any starting point only affects the TRP objective by an additive $\mathcal{O}(n)$ term.

To provide intuition on the rate $\Theta(n\sqrt{n})$, assume that the points are sampled uniformly on a compact space. For the kth served point of the TRP tour with $k \ge k^* = \lfloor n/2 \rfloor$, we have $l_k \ge l_{TSP}(k^*, n)$. Then, by Theorem 2, the expected latency of the kth point is $\Omega(\sqrt{n})$. Because this holds for all $k \ge n/2$, the expected total latency is $\Omega(n\sqrt{n})$. Similarly, we can give a simple argument for an upper bound of the expected TRP objective. Consider following the optimal TSP tour of length l_{TSP} with a starting point chosen uniformly at random among the n points. Because the position of each vertex in the tour is uniform, Equation (3) implies that the expected latency is equal to $(n-1)/2 \cdot l_{TSP} = \mathcal{O}(n\sqrt{n})$. Therefore, the expected TRP objective is $\Theta(n\sqrt{n})$ for the uniform distribution.

Let us now turn to the case of a general distribution. We show that the TRP objective is still $\Theta(n\sqrt{n})$, but we specify the dependence of the constant on the sampling distribution. We state the main asymptotic result, which we prove in the following two sections.

Theorem 3. *Assume all n vertices are drawn according to a distribution with density f on a compact space $\mathcal{K} \subset \mathbb{R}^2$. Denote by l_{TRP} the optimal TRP objective of a tour. Then,*

$$c\iint_{\mathcal{K}^2} g_f(x,y)dxdy \le \liminf_{n\to\infty} \frac{\mathbb{E}[l_{TRP}]}{n\sqrt{n}} \le \limsup_{n\to\infty} \frac{\mathbb{E}[l_{TRP}]}{n\sqrt{n}} \le C\iint_{\mathcal{K}^2} g_f(x,y)dxdy,$$

where $0 < c < C$ are two universal constants and

$$g_f(x,y) = f(y)\left(\mathbf{1}_{f(y)<f(x)} + \frac{1}{2}\cdot \mathbf{1}_{f(y)=f(x)}\right)\sqrt{f(x)}.$$

4.1. Lower Bound on the TRP

We first prove the lower bound of Theorem 3. To do so, we approximate the densities as piece-wise constant on subsquares of the compact space $\mathcal{K}$. We begin with the case of distributions on the unit square $[0,1]^2$ with piecewise-constant density of the form

$$f(x) = \sum_{1\le k\le m^2} f_k \mathbf{1}_{Q_k}(x), \tag{4}$$

where $\{Q_i\}$ is the regular partition of the unit square into m^2 subsquares of side $1/m$. Note that, because f is a density, $\sum_{k=1}^{m^2} f_k = m^2$. We denote by $f_* = \min\{f_k, f_k > 0\}$ the minimum positive density across subsquares. By construction, sampling a vertex from density f is equivalent to choosing one of the squares with a probability f_k/m^2 associated to square Q_k, then choosing a point at random uniformly in the chosen Q_k. Let $N_k = |\{i, v_i \in Q_k\}|$ denote the number of points in each subsquare. By the strong law of large numbers, we know that $N_k/n \to f_k/m^2$ almost surely.

Now, consider the optimal TRP tour. We want to restrict the problem on each of the subsquares. To do so, we can partition the tour into subpaths such that each subpath is contained completely in a subsquare Q_k (see Figure 5).

Figure 5. (Color online) Illustration of the partition procedure of a TRP tour into subpaths $\mathcal{P}_1, \ldots, \mathcal{P}_P$ corresponding to the partition of the unit square $\mathcal{K} = [0,1]^2$ into subsquares Q_k for $1 \le k \le m^2$ ($m = 3$ here). A subpath $\mathcal{P}_i$ in a subsquare Q_k that crosses completely the margin has length at least ε_m. We can then lower bound the length of that subpath in terms of number of visited vertices using Lemma 5.

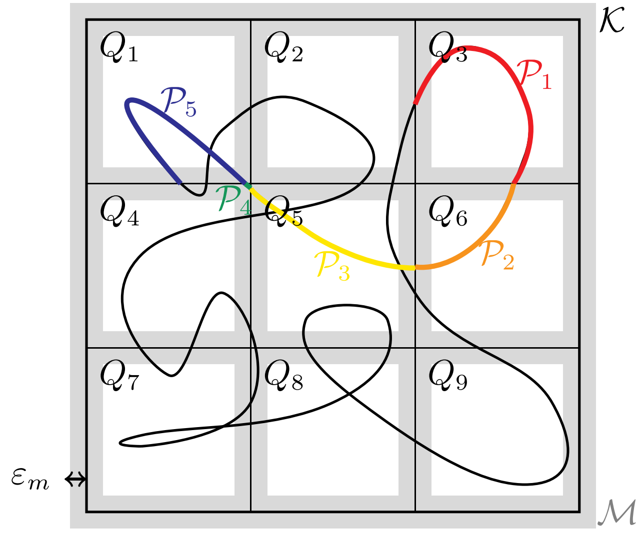

However, unlike for the TSP, we cannot "glue" the subpaths in a same subsquare Q_k directly together because here the order of subpaths impacts the TRP objective. To circumvent this issue, we derive a lower bound of the length of each subpath individually in order to obtain a lower bound on the TRP using the results on the k-TSP. To minimize the TRP objective, we order subpaths by decreasing "vertex density," defined as the ratio between the number of visited vertices in the subpath and the length of the subpath.

Define a margin $\mathcal{M}$ of the borders of the partition $\{Q_k\}$. The margin on each of the subsquares is set such that any point of Q_k outside of the margin is at a certain distance from the boundary ∂Q_k. We then are able to use Corollary 1. More precisely, denote by $B(0,1)$ the unit ball centered at the origin. Define for $\varepsilon_m := \varepsilon/m$ the margin at which $\varepsilon > 0$ is an arbitrarily small constant:

$$\mathcal{M} = \bigcup_{1\le k\le m^2} (\partial Q_k + \varepsilon_m B(0,1)).$$

Lemma 3. *We have $\mathbb{P}(|V \cap \mathcal{M}| \ge 8\varepsilon n] \le e^{-\varepsilon n}$, where $c > 0$ is a constant.*

Proof. The probability of a vertex falling inside the margin is equal to the area of the margin $\mathcal{A}_{\mathcal{M}}$. Then, $\mathcal{A}_{\mathcal{M}} \le 4m(\varepsilon/m) = 4\varepsilon$. Now, denote $c = 2/\sqrt{\pi e}$. Applying the Chernoff bound to the case of n Bernouilli $\mathcal{B}(\mathcal{A}_{\mathcal{M}})$ samples, we obtain $\mathbb{P}[|V \cap \mathcal{M}| \ge 8\varepsilon n] \le e^{-4\varepsilon n/3}$. $\square$

This lemma shows that the margin only contains a small fraction of vertices. Equivalently, most of the subpaths in Q_k visit a vertex in $Q_k \setminus \mathcal{M}$. These subpaths have length at least ε_m because they cross the margin completely. Let us now introduce the event E_0 as follows:

$$E_0 = \bigcap_{k\in\{1,\dots,m^2\}: f_k>0} \left\{ \frac{f_k}{2m^2} n \le N_k \le \frac{3f_k}{2m^2} n, \quad l_{TSP(Q_k)}\left(\left\lceil \varepsilon \cdot e\sqrt{\pi \frac{3f_k}{2m^2} n} \right\rceil, N_k \right) > \varepsilon_m \right\},$$

in which we can bound the number of points falling in each subsquare around their mean $(f_k/m^2)n$ and lower bound on the maximum number of points that can be visited by a path of length ε_m.

Lemma 4. *The event E_0 has probability $\mathbb{P}(E_0) = 1 - o\left(e^{-c\varepsilon\sqrt{(f_* n)/m^2}}\right)$ for some constant $c > 0$.*

Proof. By the Chernoff bound,

$$\mathbb{P}\left[\left| N_k - \frac{f_k}{m^2} n \right| \ge \frac{f_k}{2m^2} n \right] \le 2\exp\left(-\frac{f_k}{12m^2} n \right).$$

Moreover, using Corollary 1, we obtain for each $1 \le k \le m^2$, such that $f_k > 0$,

$$\begin{aligned}\mathbb{P}\left[l_{TSP(Q_k)}\left(\left\lceil \varepsilon \cdot e\sqrt{\pi \frac{3f_k}{2m^2} n} \right\rceil, N_k \right) \le \frac{\varepsilon}{m} \,\middle|\, \frac{f_k}{2m^2} n < N_k < \frac{3f_k}{2m^2} n \right] \\ \le \mathbb{P}\left[l_{TSP(Q_k)}\left(\left\lceil \varepsilon \cdot e\sqrt{\pi \frac{3f_k}{2m^2} n} \right\rceil, \frac{3f_k}{2m^2} n \right) \le \frac{\varepsilon}{m} \right] \\ = o\left(e^{-\varepsilon\cdot e\sqrt{(3\pi f_k n)/(2m^2)}} \right).\end{aligned}$$

Finally, we use the union bound to end the proof. $\square$

We now assume that E_0 is satisfied and analyze the length of the TRP. Recall that, in subsquare Q_k, all paths have length at least ε_m except those included in the margin $\mathcal{M}$. In particular, we can leverage the upper bound on the number of points of a path of length ε_m provided in the event E_0 to give a simple lower bound on the length of any subpath in Q_k with length at least ε_m.

Lemma 5. *Let p be a subpath in Q_k that has length $l_p \ge \varepsilon_m$ and visits n_p vertices. Then, there exists a path of length ε_m in the support of p that visits at least $\varepsilon_m n_p/(2l_p)$ vertices. Furthermore, on the event E_0, for n sufficiently large, $l_p \ge n_p(2e\sqrt{2\pi} \cdot \sqrt{f_k n})$.*

Proof. We subdivide subpath p in $\lceil l_p/\varepsilon_m \rceil$ disjoint portions of length at most ε_m. Take the portion that visits most vertices and denote by n_ε that number. In particular, $n_p \le \lceil l_p/\varepsilon_m \rceil n_\varepsilon \le (2l_p/\varepsilon_m) n_\varepsilon$, because $l_p \ge \varepsilon_m$. Note that, in

Q_k, on the event E_0, any path that visits at least $k_0 = \varepsilon \cdot e\sqrt{(3\pi f_k n)/(2m^2)}$ vertices has length at least ε_m. Therefore, $n_\varepsilon \le k_0$. Thus, for n sufficiently large, $n_p \le 2l_p(k_0/\varepsilon_m) = l_p \cdot 2e\sqrt{2\pi} \cdot \sqrt{f_k n}$. The proof follows. □

In particular, Lemma 5 shows that the vertex density of subpaths in Q_k cannot exceed the vertex density of the TSP on the N_k points in Q_k, up to a constant. We can now apply this bound to the length of all subpaths that are not completely included in the margin $\mathcal{M}$ in order to lower bound the TRP objective.

Proof of the Lower Bound of Theorem 3. First consider the case of piece-wise constant densities as defined in Equation (4). Enumerate the subpaths $\mathcal{P}_1, \ldots \mathcal{P}_P$, which are not included completely in the margin by order of appearance in the TRP path, and denote by $n(\mathcal{P}_i)$ the number of vertices visited by $\mathcal{P}_i$. Let $l(\mathcal{P}_i)$ be the length of $\mathcal{P}_i$ and $k(i)$ the index of subsquares containing $\mathcal{P}_i$, that is, $\mathcal{P}_i \subset Q_{k(i)}$. Finally, let $\tau(v)$ be the latency at point $v \in V$. On the event E_0, we can give the following lower bound on the TRP objective by applying Lemma 5 to each of the subpaths $\mathcal{P}_i$.

$$l_{TRP} = \sum_{1\le i\le P}\sum_{v\in\mathcal{P}_i}\tau(v) \ge \sum_{1\le i\le P} n(\mathcal{P}_i)\sum_{1\le j\le i-1} l(\mathcal{P}_j) \ge \frac{1}{2e\sqrt{2\pi n}}\sum_{1\le j\le P}\frac{n(\mathcal{P}_j)}{\sqrt{f_{k(j)}}}\sum_{j+1\le i\le P} n(\mathcal{P}_i). \tag{5}$$

In order to further lower bound the right term, we use the following lemma, which states that the ordering of subpaths minimizing this objective is exactly the ordering by decreasing density $f_{k(i)}$, which formalizes the intuition that it is advantageous to first serve regions with higher density.

Lemma 6. *A solution of the following minimization problem*

$$\min_{\sigma\in\mathcal{S}_P}\sum_i \frac{n(\mathcal{P}_{\sigma(i)})}{\sqrt{f_{k(\sigma(i))}}}\sum_{j>i} n(\mathcal{P}_{\sigma(j)})$$

is given by ordering the subpaths $\mathcal{P}_i$ by decreasing order of $f_{k(i)}$.

Proof. Denote by C_σ the objective of the minimization problem for $\sigma \in \mathcal{S}_P$. Let $1 \le i < j \le P$. We compare C_σ and $C_{\tilde\sigma}$, where $\tilde\sigma$ was obtained from σ by inserting the jth term in ith position. Formally, $\tilde\sigma(j) = \sigma(i)$ for $i < r \le j$, $\tilde\sigma(r) = \sigma(r-1)$, and other entries are left unchanged. Then,

$$C_{\tilde\sigma} - C_\sigma = n(\mathcal{P}_{\sigma(j)})\sum_{i\le r\le j-1} n(\mathcal{P}_{\sigma(r)})\left(\frac{1}{\sqrt{f_{k(\sigma(j))}}} - \frac{1}{\sqrt{f_{k(\sigma(r))}}}\right).$$

Assume that, for $i \le r \le j-1$, we have $1/\sqrt{f_{k(\sigma(j))}} \le 1/\sqrt{f_{k(\sigma(r))}}$. Then, the objective is decreased when we place $\sigma(j)$ in the i-th position, $C_{\tilde\sigma} \le C_\sigma$. We use this argument to order sequentially the permutation σ. First, take the index i that minimizes $1/\sqrt{f_{k(i)}}$. Let σ^* be a permutation such that $1/\sqrt{f_{k(\sigma^*(i))}}$ are in increasing order. We can first place $\sigma^*(1)$ as the first index $\tilde\sigma(1) = \sigma^*(1)$, decreasing the objective C_σ. We then place $\sigma^*(2)$ as the second index $\tilde\sigma(2) = \sigma^*(2)$ until we reach the permutation σ^* of decreasing order of $f_{k(i)}$. Thus, $C_{\sigma^*} \le C_\sigma$, and σ^* is a minimizer of the problem. □

Let us now give estimates on the right-hand side of Equation (5). Denote by σ^* the ordering on the subsquares Q_k such that $f_{\sigma^*(k)}$ is decreasing in k. Then, on the event E_0,

$$\begin{aligned}
\sum_{1\le j\le P}\frac{n(\mathcal{P}_j)}{\sqrt{f_{k(j)}}}\sum_{j+1\le i\le P} n(\mathcal{P}_i) &\ge \min_{\sigma\in\mathcal{S}_P}\sum_{i<j}\frac{n(\mathcal{P}_{\sigma(i)})}{\sqrt{f_{k(\sigma(i))}}} n(\mathcal{P}_{\sigma(j)})\\
&\ge \sum_{1\le k<t\le m^2}\frac{N_{\sigma^*(k)} - |V\cap Q_{\sigma^*(k)}\cap\mathcal{M}|}{\sqrt{f_{\sigma^*(k)}}}\cdot(N_{\sigma^*(t)} - |V\cap Q_{\sigma^*(t)}\cap\mathcal{M}|)\\
&\ge \sum_{1\le k<t\le m^2}\frac{N_{\sigma^*(k)}}{\sqrt{f_{\sigma^*(k)}}}N_{\sigma^*(t)} - \frac{2}{\sqrt{f_*}}\sum_{1\le k,t\le m^2} N_{\sigma^*(t)}|V\cap Q_{\sigma^*(k)}\cap\mathcal{M}|\\
&\ge \frac{n^2}{4m^4}\sum_{1\le k<t\le m^2}\sqrt{f_{\sigma^*(k)}}f_{\sigma^*(t)} - \frac{2n|V\cap\mathcal{M}|}{\sqrt{f_*}},
\end{aligned}$$

where, in the last inequality, we used the fact that, on E_0, $N_k \ge f_k n/(2m^2)$ for all $1 \le k \le m^2$, and $f_* = \min\{f_k : f_k > 0\}$. By Lemma 3, with probability $1 - o(\exp(-c\varepsilon\sqrt{(f_* n)/m^2}))$, the event E_0 is met and $|V \cap \mathcal{M}| \le 8\varepsilon n$. Denote by E_1 this event. Therefore, using Equation (5), on E_1,

$$l_{TRP} \ge \frac{1}{2e\sqrt{2\pi n}}\left(\frac{n^2}{4m^4}\sum_{1\le k<t\le m^2}\sqrt{f_{\sigma^*(k)}}f_{\sigma^*(t)} - \frac{16\varepsilon n^2}{\sqrt{f_*}}\right). \tag{6}$$

We now compare the right term of the preceding inequality with the integral of g_f. Note that

$$\iint_{\mathcal{K}^2} g_f(x,y)dxdy = \sum_{1\le k\le m^2} \frac{\sqrt{f_{\sigma^*(k)}}}{m^2}\int_{\mathcal{K}} f(y)\left(\mathbf{1}_{f(y)<f_{\sigma^*(k)}} + \frac{1}{2}\cdot \mathbf{1}_{f(y)=f_{\sigma^*(k)}}\right)dy$$
$$= \sum_{1\le k\le m^2} \frac{\sqrt{f_{\sigma^*(k)}}}{m^2}\left(\frac{1}{2}\frac{f_{\sigma^*(k)}}{m^2} + \sum_{k<t\le m^2}\frac{f_{\sigma^*(t)}}{m^2}\right)$$
$$= \frac{1}{2m^2}\int_{\mathcal{K}} f(x)^{3/2}dx + \frac{1}{m^4}\sum_{1\le k<t\le m^2}\sqrt{f_{\sigma^*(k)}}f_{\sigma^*(t)}.$$

The first term in the right-hand side can be made arbitrarily small. Indeed, we can repeat the complete procedure with a finest partition of the unit square $[0,1]^2$ into $(\alpha m)^2$ subsquares, where $\alpha\in\mathbb{N}^*$. For α sufficiently large, we can get

$$\frac{1}{\alpha^2 m^2}\int_{\mathcal{K}} f(x)^{3/2}dx \le \delta \iint_{\mathcal{K}^2} g_f(x,y)dxdy$$

for any arbitrarily small $\delta>0$. Then, with this partition, we have

$$\frac{1}{m^4}\sum_{1\le k<t\le m^2}\sqrt{f_{\sigma^*(k)}}f_{\sigma^*(t)} \ge (1-\delta)\iint_{\mathcal{K}^2} g_f(x,y)dxdy.$$

Therefore, taking $\varepsilon < 2^{-6}\delta\sqrt{f_*}\iint_{\mathcal{K}^2} g_f(x,y)dxdy$, Equation (6) implies that, on E_1,

$$l_{TRP} \ge \frac{1-2\delta}{8e\sqrt{2\pi}} n\sqrt{n}\iint_{\mathcal{K}^2} g_f(x,y)dxdy.$$

We now obtain the desired result,

$$\liminf_{n\to\infty}\frac{\mathbb{E}[l_{TRP}]}{n\sqrt{n}} \ge \liminf_{n\to\infty}\ \mathbb{P}[E_1]\cdot\frac{1-2\delta}{8e\sqrt{2\pi}}\iint_{\mathcal{K}^2} g_f(x,y)dxdy \ge \frac{1-2\delta}{8e\sqrt{2\pi}}\iint_{\mathcal{K}^2} g_f(x,y)dxdy.$$

This ends the proof for the densities of the form $f(x)=\sum_{k=1}^{m^2} f_k\mathbf{1}_{Q_k}(x)$. Let us now consider the general case of a distribution on a compact space $\mathcal{K}$ with both a singular part and absolutely continuous part with density f. We lower bound the TRP objective by the sum of latencies of points that do not lie in the support of the singular part. With this argument, we can restrict to the case of absolutely continuous distributions with density f without loss of generality. By a scaling argument, we can also suppose without loss of generality that $\mathcal{K}\subset[0,1]^2$. We need the following lemma to approximate f with a piece-wise constant density, the proof of which is deferred to the appendix.

Lemma 7. *Let f be a density on $\mathcal{K}\subset[0,1]^2$. For any $\varepsilon>0$, there exists a density ϕ of the form $\phi(x)=\sum_{1\le k\le m^2}\phi_k\mathbf{1}_{Q_k}(x)$ such that $\|\phi-f\|_1\le\varepsilon$ and $\left|\iint_{\mathcal{K}^2} g_\phi - \iint_{\mathcal{K}^2} g_f\right|\le\varepsilon$.*

For any $\varepsilon>0$, we use Lemma 7 to take a density ϕ of the same piece-wise constant form as in Equation (4) such that $\|\phi-f\|_1\le\varepsilon$ and $|\iint_{\mathcal{K}^2} g_\phi - \iint_{\mathcal{K}^2} g_f|\le\varepsilon$. By a coupling argument, we can construct a joint distribution (X,Y) such that X (respectively, Y) has density f (ϕ), and $\mathbb{P}(X\ne Y)\le 2\int_{\mathcal{K}}|\phi(x)-f(x)|dx\le 2\varepsilon$. Define $n_\varepsilon := |\{i, X_i\ne Y_i\}|$. Then,

$$l_{TRP(\phi)} := l_{TRP}(Y_1,\dots,Y_n) \le n[l_{TSP}(Y_i, X_i\ne Y_i)+\sqrt{2}] + l_{TRP}(Y_i, X_i=Y_i)$$
$$\le l_{TRP}(Y_i, X_i=Y_i) + 2n\sqrt{n_\varepsilon} + n(C+\sqrt{2}),$$

where in the second inequality, we use Lemma 1. Note that, using the Hoeffding inequality, we have $n_\varepsilon\le 3\varepsilon n$ with probability at least $1-e^{-2\varepsilon^2 n}$. Therefore,

$$\frac{\mathbb{E}[l_{TRP(f)}]}{n\sqrt{n}} \ge \frac{\mathbb{E}[l_{TRP}(Y_i, X_i=Y_i)]}{n\sqrt{n}} \ge \frac{\mathbb{E}[l_{TRP(\phi)}]}{n\sqrt{n}} - 2\sqrt{3\varepsilon} + o(1).$$

We can now use the result proved for density ϕ:

$$\liminf_{n\to\infty}\frac{\mathbb{E}[l_{TRP}(f)]}{n\sqrt{n}} \ge c\iint_{\mathcal{K}^2} g_\phi(x,y)dxdy - 2\sqrt{3\varepsilon} \ge c\iint_{\mathcal{K}^2} g_f(x,y)dxdy - c\cdot\varepsilon - 2\sqrt{3\varepsilon}.$$

This holds for any $\varepsilon>0$; hence, this ends the proof of the desired TRP objective upper bound. □

4.2. Upper Bound on the TRP

The proof of the lower bound of Theorem 3 from Section 4.1 suggests a procedure visiting points by zones of decreasing density. We now provide a simple construction of a tour that uses this intuition and shows the upper bound from Theorem 3.

Proof of the Upper Bound of Theorem 3. By a scaling argument, we suppose without loss of generality that $\mathcal{K} \subset [0,1]^2$. We use the same notations as in the proof of the lower bound of the expected TRP objective. Let $\varepsilon > 0$ be a tolerance parameter. Now, take $m > 0$ and a density ϕ given by Lemma 7 to approximate f. We order the subsquares by decreasing values of ϕ_k: $\phi_{\sigma(1)} \geq \cdots \geq \phi_{\sigma(m^2)}$. For each of the subsquares Q_k, we construct a tour that is optimal for the TSP; in practice, only a constant-factor approximation is needed that makes the construction polynomial: one can, for example, take the tour of Lemma 1. The output TRP tour is given by gluing together these local TSP tours into a complete tour, following the order σ. More precisely, we first follow the TSP tour in $Q_{\sigma(1)}$, and then, the TSP tour in $Q_{\sigma(2)}$ up to the TSP tour in $Q_{\sigma(m^2)}$ (see Figure 6). If a subsquare does not contain vertices, we may skip it. As a remark, the additional length for linking the subtours is negligible as $n \longrightarrow \infty$.

We now prove that this tour is constant-factor optimal with high probability. Define the event

$$E_0 = \bigcap_{1 \leq k \leq m^2} \left\{ \frac{\phi_k}{2m^2} n \leq N_k \leq \frac{3\phi_k}{2m^2} n \right\},$$

where N_k is the count of vertices in subsquare Q_k. Recall that $\mathbb{E}[N_k] = (\phi_k/m^2)n$. Therefore, using the same argument as in the proof of the lower bound, E_0 is met with probability $1 - o(\exp(-c\phi_* n/m^2))$ for some constant $c > 0$ and for which $\phi_* := \min\{\phi_k : \phi_k > 0\}$. In the next steps, we assume that E_0 is met.

By Lemma 1, if we denote by l^k_{TSP} the length of the optimal TSP tour in subsquare Q_k, then

$$l^k_{TSP} \leq \left(2\sqrt{\frac{3\phi_k}{2m^2} n} + C \right) \frac{1}{m} = \frac{\sqrt{6\phi_k n}}{m^2} + \frac{C}{m} \tag{7}$$

for all $1 \leq k \leq m^2$, and $C > 0$ a universal constant. We are now ready to estimate the TRP objective of our defined tour. Let us denote by $\hat{l}_{TRP}$ this objective and $\hat{l}_i$ the distance before visiting vertex i by following the given tour. For each subsquare Q_k, denote by i_k the index of the last vertex to be visited in this subsquare by the constructed

Figure 6. (Color online) Illustration of the constant-factor optimal TRP tour constructed for the upper bound of Theorem 3. The space is subdivided in subsquares, and the tour performs a constant-factor optimal TSP tour on each of the subsquares, following the decreasing order of density on the subsquares. The TSP tour on each subsquare is represented by a dashed path, and the density on each subsquare is represented by its brightness—darker (lighter) shades for high (low) density. Each subsquare is given a priority order from its density: the tour visits zones by decreasing order of density.

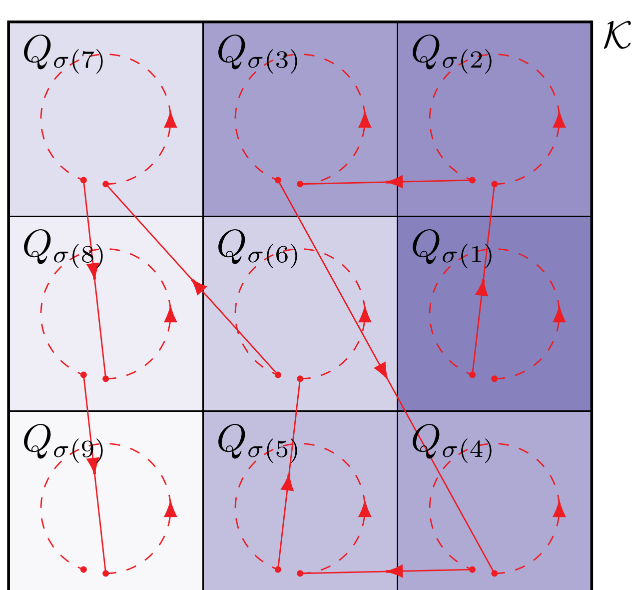

tour.

$$\hat{l}_{TRP} = \sum_{k=1}^{m^2} \sum_{i:v_i \in Q_{\sigma(k)}} \hat{l}_i \le \sum_{k=1}^{m^2} N_{\sigma(k)} \hat{l}_{i_{\sigma(k)}} \le \sum_{k=1}^{m^2} N_{\sigma(k)} \left(\sum_{l=1}^{k} l_{TSP}^{\sigma(l)} + (k-1)\sqrt{2} \right).$$

The second term $(k-1)\sqrt{2}$ was obtained by upper bounding the length of each edge linking a subsquare $Q_{\sigma(l)}$ to the next subsquare $Q_{\sigma(l+1)}$. Therefore, on E_0, because $N_k \le 3\phi_k n/(2m^2)$ for all $1 \le k \le m^2$,

$$\begin{aligned}
\hat{l}_{TRP} &\le \sum_{k=1}^{m^2} N_{\sigma(k)} \left(\sum_{l=1}^{k} l_{TSP}^{\sigma(l)} \right) + \sqrt{2}(m^2-1) \sum_{k=1}^{m^2} N_{\sigma(k)} \\
&\le \frac{3n}{2m^2} \sum_{l=1}^{m^2} l_{TSP}^{\sigma(l)} \left(\sum_{k=l}^{m^2} \phi_{\sigma(k)} \right) + \sqrt{2}(m^2-1)n \\
&\le \frac{3}{2}\sqrt{6} \frac{n\sqrt{n}}{m^4} \sum_{1 \le k \le l \le m^2} \sqrt{\phi_{\sigma(k)}} \phi_{\sigma(l)} + \frac{3C}{2m} n + \sqrt{2}(m^2-1)n,
\end{aligned}$$

where, in the last inequality, we use Equation (7). As in the proof of the lower bound of Theorem 3,

$$\frac{1}{m^4} \sum_{1 \le k \le l \le m^2} \sqrt{\phi_{\sigma(k)}} \phi_{\sigma(k)} = \iint_{\mathcal{K}^2} g_\phi(x,y) + \frac{1}{2m^2} \int_{\mathcal{K}} \phi(x)^{3/2} dx.$$

Therefore, with $\tilde{C} := 3\sqrt{6}/2$, on E_0, we obtain

$$\begin{aligned}
\frac{\hat{l}_{TRP}}{n\sqrt{n}} &\le \tilde{C} \iint_{\mathcal{K}^2} g_\phi(x,y) dx dy + \frac{\tilde{C}}{2m^2} \int_{\mathcal{K}} \phi(x)^{3/2} dx + \frac{\sqrt{2}m^2 + 3C/(2m)}{\sqrt{n}} \\
&\le \tilde{C} \iint_{\mathcal{K}^2} g_f(x,y) + \tilde{C} \left(\varepsilon + \frac{1}{2m^2} \int_{\mathcal{K}} \phi(x)^{3/2} dx \right) + \frac{\sqrt{2}m^2 + 3C/(2m)}{\sqrt{n}}.
\end{aligned}$$

Outside of the event E_0, we can use a naive upper bound $\hat{l}_{TRP} \le \sum_{i=1}^{n-1} i\sqrt{2} \le n^2/\sqrt{2}$ obtained by upper bounding the length of each edge by $\sqrt{2}$. Because $\mathbb{P}[E_0^c] = o(\exp(-c\phi_* n/m^2))$, the total contribution of this event is negligible, and we obtain

$$\limsup_{n \to \infty} \frac{\mathbb{E}[\hat{l}_{TRP}]}{n\sqrt{n}} \le \tilde{C} \iint_{\mathcal{K}^2} g_f(x,y) + \tilde{C} \left(\varepsilon + \frac{1}{2m^2} \int_{\mathcal{K}} \phi(x)^{3/2} dx \right).$$

Finally, we can take m arbitrarily large, and $\varepsilon > 0$ arbitrarily small. The result follows. □

We note that this proof of the upper bound uses an a priori algorithm to derive a TRP tour. Namely, the proposed solution visits subsquares of size $(1/m) \times (1/m)$ by decreasing order of density, using only distributional knowledge. Then, the TRP tour is adjusted by visiting the points upon realization of uncertainty by solving a TSP within each subsquare. This algorithm yields a constant-factor approximation of the optimal TRP latency. As a remark, in order for the estimates in the preceding proof to hold, we need $m^2 \ll n$ for concentration inequalities to hold on the number of points falling in each subsquare. In fact, with similar arguments, one can show that, if $\hat{l}_{TRP}$ denotes the TRP objective obtained by the procedure, whenever $m^2 \ll n$, for any $\varepsilon > 0$,

$$\mathbb{P}\left[\hat{l}_{TRP} \ge (2+\varepsilon) n\sqrt{n} \iint_{\mathcal{K}^2} g_f(x,y) dx dy \right] \underset{m \to \infty}{\longrightarrow} 0.$$

5. Fair Routing for the *k*-TSP and TRP

In the first two sections, we provide bounds for the k-TSP and TRP as well as constant-factor approximation algorithms to provide upper bounds. Both of these approximation schemes rely on a spatial discrimination approach by prioritizing the zones with high density (high probability density and high point density). Specifically, the approximation scheme for the k-TSP visits points only in the highest density zone (Section 3.3), and the approximation scheme for the TRP visits zones sequentially by decreasing order of density (Section 4.2). In fact, these schemes are derived from the lower bound analyses. This suggests that solutions to the k-TSP and TRP fundamentally integrate location-based prioritizations.

Therefore, optimizing for the k-TSP and TRP comes at the expense of spatial discrimination. In the proposed k-TSP scheme, points that do not lie on the highest density zone are never visited. Consider the simple setting in which a company can choose which customers to serve and generally receives orders from two cities. Following the proposed scheme, the company exclusively serves customers from the highest density city and, thus, ignores customers from one city altogether even though the densities might be arbitrarily close. Similarly, in the proposed TRP scheme, the waiting time is much lower in high-density than low-density regions.

To alleviate spatial discrimination outcomes, we incorporate fairness considerations into the k-TSP and TRP. Namely, we consider two categories of fairness: (i) geographic fairness, which mitigates disparities across regions, and (ii) population-based fairness, which mitigates disparities across underlying subpopulations. In the aforementioned example, under geographic fairness, the company needs to serve both cities; under population-based fairness, it needs to achieve a similar level of service across demographics (based on race or gender for instance). We quantify the efficiency–fairness trade-off via the fairness ratio, defined as the ratio between the objectives of the fair and efficient solutions. This notion relates to the price of fairness introduced by Bertsimas et al. [9], which measures the relative loss (as compared with the ratio) between the fair and efficient solutions.

5.1. Fair *k*-TSP

We focus on the case $k = o(n)$, $k \to \infty$, and points are sampled according to a continuous density f, for which Proposition 1 provides an efficient constant-factor algorithm.

5.1.1. Geographic Fairness. Denote by A_i the event in which X_i is served. By symmetry, we focus on the event A_1. A first approach to enforce fairness would be to ask that A_1 is independent of the position X_1. Stated in a more flexible way, we enforce that the probability of service conditioned on the position exceeds a threshold $\varepsilon > 0$. We define geographic fairness as follows:

$$\mathbb{P}(A_1 | X_1 = x) \geq \varepsilon \frac{k}{n}, \quad \forall x \in \mathcal{K}. \tag{8}$$

The discount factor k/n accounts for the fact that only k of the n points can be selected. Indeed, by symmetry, $\mathbb{P}(A_1) = 1/n\mathbb{E}[\mathbf{1}_{A_1} + \cdots + \mathbf{1}_{A_n}] = k/n$. The minimum service probability imposes visiting the full support of the distribution. This can be viewed as a relaxed version of max-min fairness in which we maximize the value of $\varepsilon > 0$. However, under this requirement, the k-TSP loses its locality property, inducing a significant loss in efficiency, formalized in the following proposition.

Proposition 2. *Assume that $\sqrt{n} \ll k \leq n$. Under geographic fairness (Equation (8)), the length l of a fair k-TSP path satisfies*

$$\mathbb{E}[l] \geq (1 + o_n(1))c \cdot \varepsilon \frac{k}{\sqrt{n}} \int_{\mathcal{K}} \sqrt{f},$$

where $c = 1/(e\sqrt{\pi}) > 0$ is a universal constant.

Proof. We show the result in the case of distributions on $[0,1]^2$ with piece-wise constant densities on a partition $\{Q_q\}_{q=1}^{m^2}$ defined as in Equation (4). From a given path visiting k points in the support $[0,1]^2$, we can construct a set of subpaths in each of the subsquares such that, together, they visit the same points and have same total length up to a constant dependent only on m—the length of the boundary of the subsquares partition (see Figure 7). Because $k \gg \sqrt{n}$, with high probability, $l_{TSP}(k,n) \geq ck/\sqrt{n} \to \infty$ for some constant $c > 0$. In particular, the additional constant length of the boundary is negligible compared with the length of the path visiting k points. We can now lower bound the length of the path in each subsquare separately. Denote by n_q the number of points visited by the considered path in Q_q and $B_{i,q}$ the event that X_i lies in subsquare Q_q. Under the fairness constraint (Equation (8)), we have $\mathbb{P}[A_i | B_{i,q}] \geq \varepsilon k/n$. Then, $\mathbb{E}[n_q] = n\mathbb{E}[\mathbf{1}_{A_1 \cap B_{1,q}}] \geq n\mathbb{P}[B_{1,q}] \cdot \varepsilon k/n = f_q \cdot \varepsilon k$.

Then, if l_q denotes the length of the path reduced to subsquare Q_q, we lower bound l_q with the n_q−TSP on subsquare Q_q, which has at least $(1-\eta)nf_k/m^2$ points with high probability for any fixed $\eta > 0$. Using the proof of the k-TSP lower bound (Theorem 2), with high probability, we have

$$l_q \geq l_{TSP,Q_q}(n_q, N_q) \geq \frac{n_q\sqrt{1-\eta}}{e\sqrt{\pi f_q n}} - O\left(\frac{\log^2 n}{\sqrt{n}}\right).$$

Figure 7. Consider a k-TSP path on the left figure. We partition the path into subpaths in each subsquare shown in the right figure. The length of the original k-TSP path and the sum of length of the subpaths differs at most by $\mathcal{O}(\mathcal{B})$, where $\mathcal{B}$ denotes the length of the boundary of the partition. In particular, when $k \gg \sqrt{n}$, the k-TSP length grows to infinity. Therefore, the constant boundary length is negligible compared with the k-TSP length.

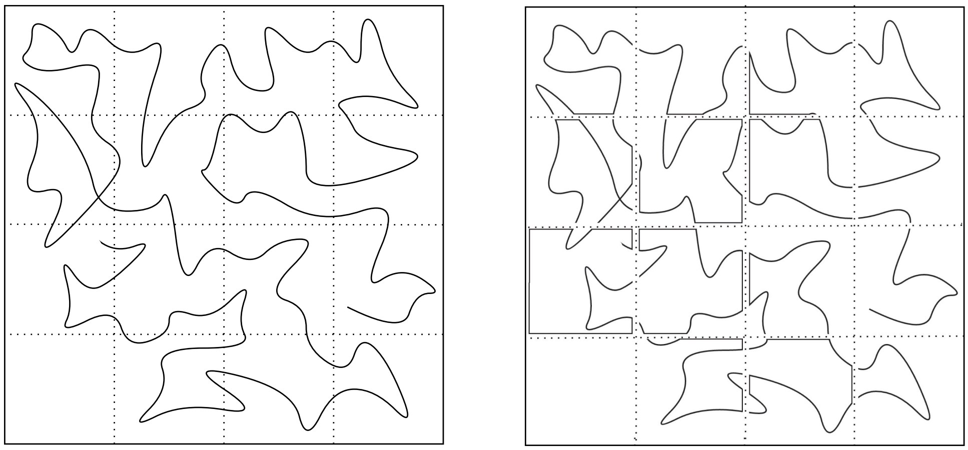

Taking the expectation and summing these inequalities yields the desired result on the fair k-TSP length, in which the $1 - o_n(1)$ term corresponds to conditioning on the high-probability event. □

As a result, the fairness ratio of a geographically fair k-TSP for $k \gg \sqrt{n}$ compared with the k-TSP length (Theorem 2) is $\Omega(\sqrt{\|f\|_\infty} \int_{\mathcal{K}} \sqrt{f})$. In Proposition 2, we assume $k \gg \sqrt{n}$ for simplicity, but the same nonlocal behavior also arises in the general case $k \to \infty$. Essentially, the geographically fair k-TSP loses the factor corresponding to the power of choosing which area to serve and the resulting fairness ratio can be arbitrarily large when the density is highly concentrated.

5.1.2. Population-Based Fairness. As suggested by the proof of probabilistic bounds in Section 3 and the fairness ratio of geographical fairness, the k-TSP is fundamentally spatially unfair. That is, the flexibility to choose which points to visit leads to disregarding zones with low density. Vice versa, imposing to visit all regions with a geographic fairness objective leads to a large loss in efficiency. In response, we now propose a second fairness notion to mitigate the price of fairness.

Consider the setting in which points belong to different populations, for instance, based on racial, gender, or age-based demographics. We aim to design solutions of the k-TSP that treat these populations fairly. For instance, one can think of a company constructing an efficient routing procedure, ensuring fairness between distinct subpopulations of customers.

Consider P populations such that points are sampled according to the density $f = f_1 + \cdots + f_P$, where f_i corresponds to the distribution of population $i = 1, \ldots, P$. For instance, we can view the sampling process as sampling a point according to density f and then assigning population i to this point with probability $f_i(X)/f(X)$. Population-based fairness asks to serve a fair number of points from each population. We propose deterministic and randomized notions of population-based fairness.

5.1.2.1. Deterministic Population-Based Fairness. A natural approach to population-based fairness involves finding a path visiting a fixed proportion p_i of points from each population $i = 1, \ldots, P$. For instance, with $p_i = 1/P$, this means that the k-TSP tour visits the same number of points from each population; with $p_i = \int f_i$, this means that the k-TSP tour serves each population proportionally to its overall size. However, we argue that this notion of fairness can be too restrictive and lead to an important loss in terms of efficiency.

Because the fair k-TSP has to visit a fixed proportion of points from each population in the same local area, we can lower bound the length of the fair k-TSP by the length of the $(p_i k)$-TSP for density f_i in this local area. In particular, the tour is constrained to visit the zone maximizing the local density of the least-represented population $\min_i f_i$, which leads to the following estimate for the length l of a fair k-TSP under deterministic population-based fairness:

$$\mathbb{E}[l] \geq (1 + o_n(1))c \cdot \frac{k}{\sqrt{\|\min_i f_i\|_\infty n}},$$

Figure 8. (Color online) Case of two populations with truncated Gaussian densities f_1 and f_2. Under deterministic population-based fairness, the fair k-TSP tour needs to visit points at which $\min(f_1, f_2)$ is maximal (i.e., at $x = 1/2$) instead of visiting points at which $f = f_1 + f_2$ is maximal ($x \approx 1/4$ or $x \approx 3/4$). In contrast, under randomized population-based fairness, the fair k-TSP tour can visit points at $x \approx 1/4$ with probability $1/2$ and points at $x \approx 3/4$ with probability $1/2$. In this example, deterministic population-based fairness yields an arbitrarily large fairness ratio and randomized population-based fairness has a fairness ratio of one.

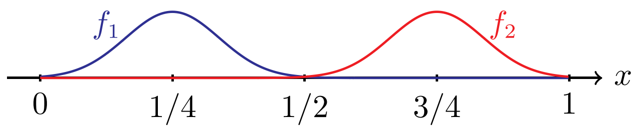

where $c > 0$ is a constant depending only on P and the fixed proportions p_i. Further, solving the k-TSP locally on the region of maximum-minimum population density $\|\min_i f_i\|_\infty$ achieves this lower bound up to a constant by Theorem 2. Hence, the efficiency fairness ratio for deterministic population-based fairness is $\Theta(\sqrt{\|f\|_\infty / \|\min_i f_i\|_\infty})$. When the populations are distributed equally over the space $\mathcal{K}$, this ratio can be close to one. In contrast, when populations are segregated, this ratio can be arbitrarily large. For instance, consider the simple case of $p = 2$ populations with truncated Gaussian densities centered in distant points. In this case, $\|\min f_i\|_\infty$ can be arbitrarily small compared with $\|f\|_\infty$ (see Figure 8 for an illustration in one dimension). Further, we can note that, if two populations do not have intersecting support, the length of any fair k-TSP is $\Omega(1)$, whereas the length of the k-TSP vanishes whenever $k \ll \sqrt{n}$. Hence, the price of fairness may still be arbitrarily large under deterministic population-based fairness.

5.1.2.2. Randomized Population-Based Fairness. In light of these limitations, randomized population-based fairness seeks a distribution of k-TSP tours as opposed to a single solution. We only ensure that the k-TSP tour visits a fixed proportion p_i of points from each population in expectation, but every single k-TSP tour may deviate from the proportions p_i. Again, $p_i = 1/P$ corresponds to equal service (in expectation), and $p_i = \int f_i$ corresponds to proportional service (in expectation). Such randomization allows for more flexibility than deterministic fairness because individual paths of the output distribution can possibly serve populations heterogeneously.

For simplicity, we consider the case in which densities f_i are piece-wise constant on a partition $\{Q_j\}$ of the unit square $[0,1]^2$ in m^2 subsquares of equal size $(1/m) \times (1/m)$, that is, $f_i = \sum_{j=1}^{m^2} f_{i,j} \mathbf{1}_{Q_j}$. We can relax this assumption by approximating continuous densities with piece-wise densities on the partition for large enough m. However, this simplification is useful to provide intuition on the proposed randomized population-based fairness scheme. We write the total density as $f = \sum_{j=1}^{m^2} f_j \mathbf{1}_{Q_j}$. Without loss of generality, we can omit subsquares that do not contain points and assume that the total density is positive $f_j > 0$ for all subsquares Q_j. Recall that the condition $k \ll n$ ensures that a path visiting k points can be constructed locally for n large enough. We analyze the randomized approximating scheme in which we select a subsquare Q_j with probability q_j and then compute an approximating k-TSP path in this subsquare, using the algorithm proposed in Section 3.2. By symmetry, the k-TSP in subsquare Q_j visits $(f_{i,j}/f_j)k$ points from population j in expectation. Therefore, the randomized fairness constraint for our scheme imposes that

$$\sum_{j=1}^{m^2} q_j \frac{f_{i,j}}{f_j} = p_i, \quad \forall i = 1, \ldots, P. \tag{9}$$

By Theorem 2, if l denotes the length of the k-TSP path output by the randomized scheme,

$$\mathbb{E}[l] = \Theta\left(\sum_{j=1}^{m^2} q_j \frac{k-1}{(f_j n)^{\frac{1}{2}\left(1+\frac{1}{k-1}\right)}} \right), \tag{10}$$

where $c > 0$ is a universal constant. The optimal set of probabilities q_j can be obtained by solving a simple linear program minimizing the objective (Equation (10)) under the population-based fairness constraint (Equation (9)) on the probability simplex. We obtain

$$\begin{aligned}
&\min \sum_{j=1}^{m^2} q_j f_j^{-\frac{1}{2}\left(1+\frac{1}{k-1}\right)},\\
&\text{s.t.} \sum_{j=1}^{m^2} q_j \frac{f_{i,j}}{f_j} = p_i, \qquad \forall i = 1, \dots, P,\\
&\sum_{j=1}^{m^2} q_j = 1,\\
&q_i \geq 0, \qquad\qquad \forall i = 1, \dots, P.
\end{aligned}$$

Summing all fairness constraints (Equation (9)) shows that the preceding linear program contains at most P linearly independent equations. Thus, there exist an optimal probability q^* with at most P positive entries. In other words, instead of visiting all m^2 subsquares, there exists an optimal strategy for the randomized fair scheme visiting at most P different subsquares.

For instance, consider completely segregated populations, that is, populations with disjoint support. Recall that, in this setting, deterministic population-based fairness has an infinite fairness ratio. This is not the case for randomized population-based fairness. Specifically, under randomized population-based fairness, an optimal strategy consists of choosing one subsquare that maximizes the density f_i for each population, then randomly selecting the subsquare to perform the k-TSP consistently with the fairness constraints (see Figure 8 for an illustration in one dimension).

We can also add a tolerance $\varepsilon \geq 0$ for the fairness by relaxing Equation (9) to

$$p_i - \varepsilon \leq \sum_{j=1}^{m^2} q_j \frac{f_{i,j}}{f_j} \leq p_i + \varepsilon, \quad \forall i = 1, \dots, P.$$

This constraint yields a new linear program for which there still exists an optimal sparse solution q^* with at most $P+1$ nonzero entries. The ε tolerance acts as a regularization term. When $\varepsilon \geq 1$, the corresponding algorithm is blind to the fairness constraints, thus amounting to the k-TSP in the case $k \gg 1$: it only visits points in the maximum density subsquare (Section 3.3). On the other hand, when $\varepsilon = 0$, we recover the strict fairness constraint Equation (9). Denoting by q^ε the optimal probability distribution, the fairness ratio when $1 \ll k \ll n$ corresponds to the ratio of the linear program objective for the chosen tolerance parameter ε and the objective for tolerance one, that is, $\Theta\left(\sqrt{\max_j f_j} \sum_{j=1}^{m^2} q_j^\varepsilon / \sqrt{f_j}\right)$, strictly improving over the fairness ratio for deterministic fairness.

5.2. Fair TRP

Recall from Section 5.1.1 that geographical fairness can result in a significant loss in the objective of the for the k-TSP. This stems from the fact that the k-TSP is fundamentally local for small k (e.g., $k \ll \sqrt{n}$). In contrast, the TRP has a global objective and visits all points in the space. In this section, we see that our approximation scheme for the TRP can be adapted to geographical fairness without loss in fairness ratio, in particular, under max-min fairness. Additional results for other utility-based notions of fairness are given in the companion report (Blanchard et al. [13]), in which we show that the approximation scheme for the TRP can be efficiently adapted to account for this notion of fairness. In the game-theoretical setting, max-min fairness yields a Pareto optimal allocation by maximizing the minimum utility that all players derive (Bertsimas et al. [9]). In particular, whenever there exist efficient allocations in which all players have the same utility, max-min fairness outputs this equitable allocation. In the case of the TRP, we model the utility of a point by a decreasing function of its latency. In this case, max-min fairness seeks the tour visiting all n points and minimizing the worst latency, that is, the latency of the point that is visited last. In other words, max-min fairness is equivalent to the TSP, which minimizes the total tour length. We show that our proposed algorithm for the TRP in Section 4.2 is asymptotically optimal for the TSP, hence max-min fair.

Proposition 3. *The approximation algorithm for the TRP described in Section 4.2 is asymptotically max-min fair. Specifically, let l(TRP) be the maximum point latency for a TRP tour and l^* be the minimum maximum point-latency, that is, the maximum point-latency of a max-min fair allocation. Then,* $\mathbb{E}[l(TRP)] = (1 + o_n(1))\mathbb{E}[l^*]$.

Proof. The approximation algorithm for the TRP consists in serving subsquares sequentially by order of decreasing density. If σ denotes this ordering, we first perform the TSP on subsquare $Q_{\sigma(1)}$ and then on $Q_{\sigma(2)}$ until $Q_{\sigma(m^2)}$. Note that the total length of the edges linking subsquares is at most $\mathcal{O}(m^2) = \mathcal{O}(1)$, which is negligible compared with the total length of the tour $\Theta(\sqrt{n})$. We can then apply the BHH theorem to each subsquare to obtain

$$\mathbb{E}[l(TRP)] = (1 + o_n(1))\beta_{TSP}\sqrt{n}\int_{\mathcal{K}}\sqrt{f} + \mathcal{O}(1) = (1 + o_n(1))\mathbb{E}[l(TSP)],$$

where $l(TSP) = l^*$ is the length of the optimal TSP tour (hence, a max-min fair tour). □

6. Conclusion

In this paper, we give constant-factor probabilistic estimates for the k-TSP and TRP when points are sampled independently according to a known distribution. Specifically, we show that the optimal k-TSP tour grows at a rate of $\Theta(k/n^{k/(2(k-1))})$ and the optimal TRP latency grows at a rate of $\Theta(n\sqrt{n})$. Moreover, our proofs for the upper bounds are constructive based on intuitive approximation schemes. For the k-TSP, a constant-factor approximation algorithm involves performing a TSP tour in a zone with high point concentration. For the TRP, a constant-factor approximation algorithm involves creating a master a priori tour by visiting zones of decreasing probability density and then performing a TSP tour within each zone. We also proposed adaptations of these algorithms to capture fairness considerations, namely, randomized population-based fairness for the k-TSP and geographic fairness for the TRP. As discussed in Section 2.3, these results can have significant practical implications for the design of transportation and logistics systems in which the operator strives to minimize customer or passenger wait times as opposed to merely minimizing operating costs or travel times.

It is worth noting that we analyze the k-TSP and TRP in the Euclidean plane, but the results can be generalized to Euclidean spaces of higher dimension with additional technicality. Furthermore, the upper bound given for the TRP uses the master-tour construction from Lemma 1 in order to approximate the TSP locally, which yields a simple a priori algorithm. However, directly using the TSP as subroutine improves the constant two in the upper bound for the TRP to β_{TSP}, the constant appearing in the asymptotic length of the TSP. A natural question is whether this constant β_{TSP} is tight. This gives an equivalence result of the TRP latency as opposed to our constant-factor estimates. However, in our analysis, improving the constant of our lower bound for the TRP requires improving the constant of the k-TSP lower bound. In particular, this asks whether, for large k (e.g., $k = \Omega(\log n)$), the length of the k-TSP is $\sim \beta_{TSP}k/\sqrt{n}$. We leave this question open for future research. Finally, we refer to the companion report (Blanchard et al. [13]) for additional extensions on the k-TSP bounds and the fair TRP.

Acknowledgments

The authors thank Bart van Parys for valuable feedback on the manuscript.

Appendix. Proof of Lemma 7

Lemma 7 (Repeated for Convenience). *Let f be a density on $\mathcal{K} \subset [0,1]^2$. For any $\varepsilon > 0$, there exists a density ϕ of the form $\phi(x) = \sum_{1 \le k \le m^2}\phi_k \mathbf{1}_{Q_k}(x)$ such that $\|\phi - f\|_1 \le \varepsilon$ and $\left|\iint_{\mathcal{K}^2} g_\phi - \iint_{\mathcal{K}^2} g_f\right| \le \varepsilon$.*

Proof. By the Cauchy–Schwarz inequality, $\|\sqrt{f}\|_1 \le \sqrt{\|f\|_1} = 1$. Let $\varepsilon > 0$ and M such that $\int_{\mathcal{K}} f\,\mathbf{1}_{f>M} \le \varepsilon$. Then, we can take a density ϕ_ε of the right form such that $\|\phi_\varepsilon - f\|_1 \le \varepsilon, \varepsilon^{3/2}/M$ and $\|\sqrt{\phi_\varepsilon} - \sqrt{f}\|_1 \le \varepsilon$. We can also choose ϕ_ε such that all ϕ_k are distinct. For the sake of simplicity, we write ϕ instead of ϕ_ε for the next derivations. Again, we have $\|\sqrt{\phi}\|_1 \le 1$. First,

$$\iint_{\mathcal{K}^2} g_f(\mathbf{1}_{|\phi(x)-f(x)|\ge\sqrt{\varepsilon}} + \mathbf{1}_{|\phi(y)-f(y)|\ge\sqrt{\varepsilon}})dxdy \le \int_{\mathcal{K}}\sqrt{f(x)}\mathbf{1}_{|\phi(x)-f(x)|\ge\sqrt{\varepsilon}}dx + \int_{\mathcal{K}} f(y)\mathbf{1}_{|\phi(y)-f(y)|\ge\sqrt{\varepsilon}}dy.$$

By Cauchy–Schwarz, $\int_{\mathcal{K}}\sqrt{f(x)}\mathbf{1}_{|\phi(x)-f(x)|\ge\sqrt{\varepsilon}}dx \le \sqrt{\int_{\mathcal{K}}\mathbf{1}_{|\phi(x)-f(x)|>\sqrt{\varepsilon}}dx} \le \sqrt{\|\phi - f\|_1/\sqrt{\varepsilon}} \le \varepsilon^{1/4}$, where we use Markov's inequality. Also,

$$\int_{\mathcal{K}} f(y)\mathbf{1}_{|\phi(y)-f(y)|\ge\sqrt{\varepsilon}}dy \le \int_{\mathcal{K}} f(y)\mathbf{1}_{f(y)>M} + M\int_{\mathcal{K}}\mathbf{1}_{|\phi(y)-f(y)|\ge\sqrt{\varepsilon}}dy \le \varepsilon + M\frac{\|\phi - f\|_1}{\sqrt{\varepsilon}} \le 2\varepsilon.$$

Similarly, we obtain

$$\iint_{\mathcal{K}^2} g_\phi(\mathbf{1}_{|\phi(x)-f(x)|\geq\sqrt{\varepsilon}} + \mathbf{1}_{|\phi(y)-f(y)|\geq\sqrt{\varepsilon}})dxdy \leq \varepsilon^{1/4} + \int_{\mathcal{K}} \phi(y)\mathbf{1}_{|\phi(y)-f(y)|\geq\sqrt{\varepsilon}}dy$$
$$\leq \varepsilon^{1/4} + \|\phi - f\|_1 + \int_{\mathcal{K}} f(y)\mathbf{1}_{|\phi(y)-f(y)|\geq\sqrt{\varepsilon}}dy$$
$$\leq \varepsilon^{1/4} + 3\varepsilon.$$

It now remains to bound the integral of $g_f - g_\phi$ when $|\phi(x)-f(x)|, |\phi(y)-f(y)| < \sqrt{\varepsilon}$.

$$\left|\iint_{\mathcal{K}^2} (g_f - g_\phi)\mathbf{1}_{|\phi(x)-f(x)|,\,|\phi(y)-f(y)|<\sqrt{\varepsilon}}dxdy\right|$$
$$\leq \iint_{\mathcal{K}^2}\left|f(y)\sqrt{f(x)} - \phi(y)\sqrt{\phi(x)}\right|dxdy$$
$$+\left|\iint_{\substack{x,y\in\mathcal{K}\\ |\phi(x)-f(x)|,\,|\phi(y)-f(y)|<\sqrt{\varepsilon}}}\left(\mathbf{1}_{\phi(y)<\phi(x)} + \frac{\mathbf{1}_{\phi(y)=\phi(x)}}{2} - \mathbf{1}_{f(y)<f(x)} - \frac{\mathbf{1}_{f(y)=f(x)}}{2}\right)f(y)\sqrt{f(x)}dxdy\right|$$
$$\leq \iint_{\mathcal{K}^2} f(y)\left|\sqrt{f(x)} - \sqrt{\phi(x)}\right|dxdy + \iint_{\mathcal{K}^2}\left|f(y)-\phi(y)\right|\sqrt{\phi(x)}dxdy$$
$$+\left|\iint_{x,y\in\mathcal{K},\,|f(x)-f(y)|<2\sqrt{\varepsilon}}\left(\mathbf{1}_{\phi(y)<\phi(x)} + \frac{\mathbf{1}_{\phi(y)=\phi(x)}}{2} - \mathbf{1}_{f(y)<f(x)} - \frac{\mathbf{1}_{f(y)=f(x)}}{2}\right)f(y)\sqrt{f(x)}dxdy\right|$$
$$\leq 2\varepsilon + \frac{1}{2m^2}\|\phi\|_\infty^{3/2} + \left|\iint_{x,y\in\mathcal{K},\,|f(x)-f(y)|<2\sqrt{\varepsilon}}\left(\mathbf{1}_{\phi(y)<\phi(x)} - \mathbf{1}_{f(y)<f(x)} - \frac{\mathbf{1}_{f(y)=f(x)}}{2}\right)f(y)\sqrt{f(x)}dxdy\right|.$$

Now, consider the function $g(z) := \int_{\mathcal{K}}\mathbf{1}_{f(x)=z}dx$. Note that $0 \leq g \leq 1$ and $\sum_{z\geq 0} g(z) \leq 1$. Therefore, the support of g is countable $Supp(g) = \{z_i;\ i \geq 1\}$. Then, $\iint_{\mathcal{K}^2}\mathbf{1}_{f(x)=f(y)}f(y)\sqrt{f(x)}dxdy = \sum_i z_i^{3/2}g(z_i)^2$. We now look at the other terms. First note that

$$\iint_{f(x)=f(y)=z_i}\mathbf{1}_{\phi(y)<\phi(x)}dxdy = \frac{1}{2}\left(\iint_{f(x)=f(y)=z_i}\mathbf{1}_{\phi(y)<\phi(x)}dxdy + \iint_{f(x)=f(y)=z_i}\mathbf{1}_{\phi(y)>\phi(x)}dxdy\right)$$
$$= \frac{g(z_i)^2}{2} - \frac{1}{2}\iint_{f(x)=f(y)=z_i}\mathbf{1}_{\phi(y)=\phi(x)}dxdy.$$

Therefore,

$$\left|\iint_{f(x)=f(y)}\mathbf{1}_{\phi(y)<\phi(x)}f(y)\sqrt{f(x)}dxdy - \frac{1}{2}\sum_i z_i^{3/2}g(z_i)^2\right| \leq \sum_i\left|\iint_{f(x)=f(y)=z_i}\mathbf{1}_{\phi(y)<\phi(x)}f(y)\sqrt{f(x)}dxdy - \frac{1}{2}z_i^{3/2}g(z_i)^2\right|$$
$$= \sum_i z_i^{3/2}\left|\iint_{f(x)=f(y)=z_i}\mathbf{1}_{\phi(y)<\phi(x)}dxdy - \frac{g(z_i)^2}{2}\right| \leq \frac{1}{2}\iint_{\phi(y)=\phi(x)} f(x)\sqrt{f(y)}dxdy.$$

Because ϕ_k is distinct on each subsquare Q_k, by the dominated convergence theorem, the right term vanishes when m grows. Indeed, $\mathbf{1}_{\phi(y)=\phi(x)} \to \mathbf{1}_{x=y}$ as $m\to\infty$, and $\{x=y\}$ is a negligible set. Now, we take m sufficiently large such that the right term is upper bounded by δ. Finally,

$$\left|\iint_{\mathcal{K}^2} g_f - g_\phi\right| \leq 2\varepsilon^{1/4} + 7\varepsilon + \frac{1}{2m^2}\|\phi\|_\infty^{3/2} + \delta + \iint_{\mathcal{K}^2}\mathbf{1}_{|f(x)-f(y)|<2\sqrt{\varepsilon}}\mathbf{1}_{f(x)\neq f(y)}f(x)\sqrt{f(y)}dxdy.$$

By the dominated convergence theorem, the right term vanishes as $\varepsilon \to 0$. Then, taking $0 \leq \varepsilon \leq \delta$ sufficiently small and then m sufficiently large, we can achieve $|\iint_{\mathcal{K}^2} g_f - g_\phi| \leq 2\delta$. Note that we also have $\|\phi_\varepsilon - f\|_1 \leq \varepsilon \leq \delta$. This ends the proof of the lemma. □

References

[1] Afrati F, Cosmadakis S, Papadimitriou CH, Papageorgiou G, Papakostantinou N (1986) The complexity of the travelling repairman problem. *RAIRO Theoretical Informatics Appl.* 20(1):79–87.
[2] Arora S, Karakostas G (2006) A $2+\varepsilon$ approximation algorithm for the k-MST problem. *Math. Programming* 107(3):491–504.

[3] Arya S, Ramesh H (1998) A 2.5-factor approximation algorithm for the *k*-MST problem. *Inform. Processing Lett.* 65(3):117–118.
[4] Ausiello G, Leonardi S, Marchetti-Spaccamela A (2000) On salesmen, repairmen, spiders, and other traveling agents. Bongiovanni G, Gambosi G, Petreschi R, eds. *Italian Conf. Algorithms Complexity* (Springer), 1–16.
[5] Banerjee D, Erera AL, Toriello A (2023) Fleet sizing and service region partitioning for same-day delivery systems. *Transportation Sci.* 56(5):1327–1347.
[6] Banerjee D, Erera AL, Stroh AM, Toriello A (2022) Who has access to e-commerce and when? Time-varying service regions in same-day delivery. *Transportation Res. Part B: Methodological* 170:148–168.
[7] Beardwood J, Halton JH, Hammersley JM (1959) The shortest path through many points. *Math. Proc. Cambridge Philos. Soc.* 55(4):299–327.
[8] Bertsimas D, Gupta S (2016) Fairness and collaboration in network air traffic flow management: An optimization approach. *Transportation Sci.* 50(1):57–76.
[9] Bertsimas D, Farias V, Trichakis N (2011) The price of fairness. *Oper. Res.* 59(1):17–31.
[10] Bertsimas D, Farias VF, Trichakis N (2012) On the efficiency-fairness trade-off. *Management Sci.* 58(12):2234–2250.
[11] Bertsimas DJ, Jaillet P, Odoni AR (1990) A priori optimization. *Oper. Res.* 38(6):1019–1033.
[12] Bianco L, Mingozzi A, Ricciardelli S (1993) The traveling salesman problem with cumulative costs. *Networks* 23(2):81–91.
[13] Blanchard M, Jacquillat A, Jaillet P (2022) Additional results and extensions for the paper "Probabilistic bounds on the *k*-traveling salesman problem and the traveling repairman problem". Preprint, submitted November 20, https://arxiv.org/abs/2211.11065.
[14] Blum A, Ravi R, Vempala S (1996) A constant-factor approximation algorithm for the *k*-MST problem. Miller GL, ed. *Proc. 28th Annual ACM Sympos. Theory Comput.* (Association for Computing Machinery, New York), 442–448.
[15] Blum A, Chalasani P, Coppersmith D, Pulleyblank B, Raghavan P, Sudan M (1994) The minimum latency problem. Leighton T, Goodrich M, eds. *Proc. 26th Annual ACM Sympos. Theory Comput.*, 163–171.
[16] Carlsson JG (2012) Dividing a territory among several vehicles. *INFORMS J. Comput.* 24(4):565–577.
[17] Carlsson JG, Jones B (2022) Continuous approximation formulas for location problems. *Networks* 80(4):407–430.
[18] Carlsson JG, Song S (2018) Coordinated logistics with a truck and a drone. *Management Sci.* 64(9):4052–4069.
[19] Chaudhuri K, Godfrey B, Rao S, Talwar K (2003) Paths, trees, and minimum latency tours. Sudan M, eds. *44th Annual IEEE Sympos. Foundations Comput. Sci. Proc.* (IEEE, Piscataway, NJ), 36–45.
[20] Fischetti M, Hamacher HW, Jørnsten K, Maffioli F (1994) Weighted k-cardinality trees: Complexity and polyhedral structure. *Networks* 24(1):11–21.
[21] Garg N (1996) A 3-approximation for the minimum tree spanning *k* vertices. Raghavan P, eds. *Proc. 37th Conf. Foundations Comput. Sci.* (IEEE, Piscataway, NJ), 302–309.
[22] Garg N (2005) Saving an epsilon: A 2-approximation for the *k*-MST problem in graphs. Gabow H, Fafin R, eds. *Proc. 37th Annual ACM Sympos. Theory Comput.* (Association for Computing Machinery, New York), 396–402.
[23] Goemans M, Kleinberg J (1998) An improved approximation ratio for the minimum latency problem. *Math. Programming* 82(1):111–124.
[24] Jacquillat A, Vaze V (2018) Interairline equity in airport scheduling interventions. *Transportation Sci.* 52(4):941–964.
[25] Jaillet P (1993) Analysis of probabilistic combinatorial optimization problems in Euclidean spaces. *Math. Oper. Res.* 18(1):51–70.
[26] Johnson DS, Minkoff M, Phillips S (2000) The prize collecting Steiner tree problem: Theory and practice. Shmoys D, ed. *Proc. 11th Annual ACM-SIAM Sympos. Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia), 760–769.
[27] Laporte G, Louveaux FV, Mercure H (1994) A priori optimization of the probabilistic traveling salesman problem. *Oper. Res.* 42(3):543–549.
[28] Luo H, Lu S, Bharghavan V, Cheng J, Zhong G (2004) A packet scheduling approach to QOS support in multihop wireless networks. *Mobile Networks Appl.* 9(3):193–206.
[29] Mallows C (1968) An inequality involving multinomial probabilities. *Biometrika* 55(2):422–424.
[30] Minieka E (1989) The delivery man problem on a tree network. *Ann. Oper. Res.* 18(1):261–266.
[31] Navidi F, Gørtz IL, Nagarajan V (2020) Approximation algorithms for the a priori traveling repairman. *Oper. Res. Lett.* 48(5):599–606.
[32] O'Cinneide C, Scherer B, Xu X (2006) Pooling trades in a quantitative investment process. *J. Portfolio Management* 32(4):33–43.
[33] Pandiri V, Singh A (2020) Two multi-start heuristics for the k-traveling salesman problem. *Opsearch* 57(4):1164–1204.
[34] Paul A, Freund D, Ferber A, Shmoys DB, Williamson DP (2020) Budgeted prize-collecting traveling salesman and minimum spanning tree problems. *Math. Oper. Res.* 45(2):576–590.
[35] Picard JC, Queyranne M (1978) The time-dependent traveling salesman problem and its application to the tardiness problem in one-machine scheduling. *Oper. Res.* 26(1):86–110.
[36] Radunovic B, Le Boudec JY (2007) A unified framework for max-min and min-max fairness with applications. *IEEE/ACM Trans. Networking* 15(5):1073–1083.
[37] Ravi R, Sundaram R, Marathe MV, Rosenkrantz DJ, Ravi SS (1996) Spanning trees—Short or small. *SIAM J. Discrete Math.* 9(2):178–200.
[38] Simchi-Levi D, Berman O (1991) Minimizing the total flow time of *n* jobs on a network. *IIE Trans.* 23(3):236–244.
[39] Sitters R (2002) The minimum latency problem is NP-hard for weighted trees. Cook WJ, Schulz AS, eds. *Internat. Conf. Integer Programming Combin. Optim.* (Springer, Berlin), 230–239.
[40] Sitters R (2014) Polynomial time approximation schemes for the traveling repairman and other minimum latency problems. Chekuri C, ed. *Proc. 25th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 604–616.
[41] Steele JM (1981) Subadditive Euclidean functionals and nonlinear growth in geometric probability. *Ann. Probab.* 9(3):365–376.
[42] Steele JM (1997) *Probability Theory and Combinatorial Optimization* (SIAM, Philadelphia).
[43] Stroh AM, Erera AL, Toriello A (2022) Tactical design of same-day delivery systems. *Management Sci.* 68(5):3444–3463.
[44] Tsitsiklis JN (1992) Special cases of traveling salesman and repairman problems with time windows. *Networks* 22(3):263–282.
[45] van Ee M, Sitters R (2018) The a priori traveling repairman problem. *Algorithmica* 80(10):2818–2833.
[46] Vossen T, Ball M, Hoffman R, Wambsganss M (2003) A general approach to equity in traffic flow management and its application to mitigating exemption bias in ground delay programs. *Air Traffic Control Quart.* 11(4):277–292.