Coordinated logistics with a truck and multiple sidekicks

Bo Jones, Ke Xu, and John Gunnar Carlsson

Abstract

One of the more novel recent innovations in the logistics world, both in theory and in practice, is the use of small autonomous vehicles to facilitate last-mile delivery. One particular scheme that has received considerable recent attention is the "sidekick" scheme, in which a large cargo truck acts as a mobile "host" that deploys smaller vehicles, such as aerial drones or unmanned ground vehicles (UGVs). In this paper, we develop a continuous approximation model that estimates the improvements to total completion time that such a system provides, in the asymptotic limit as many demand points are drawn from a continuous probability distribution in the plane. Our key finding is that sidekick systems can be beneficial even when the sidekicks are slower than the host, provided there are sufficiently many of them.

1 Introduction

One logistical paradigm that has received considerable attention in recent years is the *sidekick routing scheme*. A sidekick routing scheme is a logistical framework in which a large "host" vehicle, such as a truck or van, serves as a mobile base for a fleet of small vehicles (the "sidekicks"), such as unmanned ground vehicles (UGVs) or unmanned aerial vehicles (UAVs). The sidekicks alternate between visiting the truck to pick up items and visiting the customers, and the overall objective is to determine a coordinated set of routes for all vehicles in order to optimize system efficiency, such as minimizing the time to completion, the vehicle miles travelled (VMT), or some other measure. A sketch of such a system is shown in Figure 1. The same model applies if we think of the sidekicks as picking up items from the customers, but for consistent, brief terminology we will place ourselves in the delivery setting throughout this work.

Until recently, the use of sidekick routing schemes was restricted to conceptual prototypes, such as the Amazon patent [10] and a pilot project by UPS [27]. However, the FAA Reauthorization Act of 2024 and subsequent regulatory progress toward standardized Beyond Visual Line of Sight (BVLOS) operations have created a more favorable climate for commercial drone delivery, with Transportation Secretary Duffy signaling imminent rules to expand drone deliveries [1] and Canada amending its Aviation Regulations to ease BVLOS restrictions for certain operations beginning November 2025 [29]. In the context of material handling in a warehouse, 6 River Systems has deployed a system called "Chuck", in which human order pickers and AGVs work collaboratively, although the problem attributes are somewhat distinct from those studied in this paper [22]. The issues in deploying sidekick systems appear to stem from a combination of regulatory hurdles as well as the practical difficulty in having physical coordinated interactions between different transportation modes; for example, the

Chuck system coordinates humans with AGVs in a way that requires fine-tuned interaction beyond the scope of existing autonomous sytems. At present, current mainstream drone delivery operations focus more on centralized systems in which drones fly directly from a central depot or warehouse to individual customers [24], i.e. without being deployed from a mobile truck. These non-sidekick systems feature simpler logistics and do not require coordination between a truck and drones. While there has been extensive academic research on truck-drone routing problems in recent years, as we discuss in our literature review, the practical implementation of these systems has been more limited, with hardware development outpacing operational deployment. From the perspective of routing these systems pose an exceptionally difficult challenge due to the need to synchronize multiple vehicles that can all be traveling at the same time and at different speeds. We cannot consider vehicles' routes separately as we must include the possibility of vehicles carrying other vehicles for periods of time and the need for intermittent meetings of vehicles at the same position at the same point in time. Thus we see that individual vehicles' routes are highly interdependent, and any reasonable objective will be impacted by this interdependence, making the optimization very hard. Furthermore, the high-level attributes of these systems are not at all clear: how much more efficient can they be? When are they useful? What are the trade-offs inherent in such a scheme? We employ a continuous approximation analysis as a means of helping to answer these questions.

This paper is organized as follows. In Section 2 we provide an overview of related work. In Section 3 we formally define the sidekick routing problem. In Section 4 we introduce preliminary results that will be of use in our analysis. In Section 5 we derive our main result concerning the asymptotic behavior of the sidekick routing problem; our key finding is that sidekick systems can be beneficial even when the sidekicks are slower than the host, provided there are sufficiently many of them. In Section 6 we summarize the operational implications of this result. In particular, we consider how much improvement in efficiency can be gained by switching to the sidekick system and the tradeoffs that must be weighed in implementing the system. Finally, in Section 7 we consider the scaling behavior, and dependence on the configuration of the sidekicks, that our result tells us we should expect. We empirically demonstrate that actual tour times, obtained by heuristically solving the sidekick problem on simulated sets of customer points, corroborate our expectations.

1.1 Remark on notational conventions

When it is necessary to specify that a map is injective, we use $\hookrightarrow$. For optimization problems we use sans serif problem names to denote the optimal objective value of the problem, e.g. the length of the optimal TSP tour is denoted TSP. We use the notation $a \propto b$ to denote linear proportionality between a and b with an unspecified proportionality constant.

2 Related work

In 2018 Otto et al. compiled a comprehensive review, [19], of work on optimization approaches to systems employing drones for a wide range of applications, to include package delivery and specifically package delivery using drones as sidekicks for trucks. More recent surveys have focused exclusively on

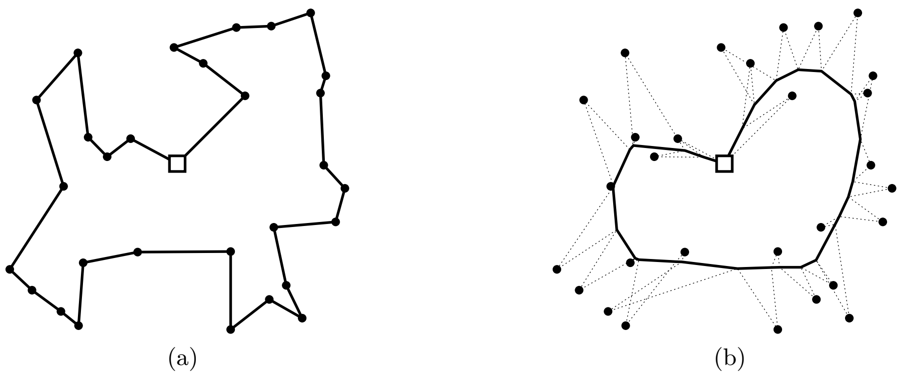

Figure 1: The figure on the left shows a *travelling salesman tour* of a set of client destinations and a central depot, that is, the shortest tour that visits a collection of points and starts and ends at the central depot. The figure on the right shows a solution to a "sidekick" problem in which the truck has a sidekick (such as a robot or a drone) that alternates between visiting the truck and visiting the customer locations with the goal of minimizing the completion time of the tour.

the routing problem for drone-aided systems; see Khoufi et al. [13] from 2019 and Macrina et. al. [15] from 2020. Viloria et al. [23] from 2020 and Moshref-Javadi and Winkenbach [17] from 2021 surveyed work on a broader class of routing problems with drones to include sidekick routing. Chung et al.'s [9] review from 2020 covers optimization problems for truck-drone coordinated systems to include delivery.

We are primarily interested in prior work in the area of continuous approximation and theoretical results bounding the objective or characterizing the improvement due to sidekick introduction. For work on solving sidekick problems we concern ourselves principally with these papers' formulations of the problem and focus on papers that contribute new model elements. There are many variants, each differing in the assumptions that are made about the delivery system. Three critical questions, the answers to which change from model to model, that need to be posed are given below.

- Does the truck also deliver packages or are packages only delivered by the sidekicks?
- Can the truck carry multiple sidekicks capable of making simultaneous deliveries?
- Are the sidekick launch and pickup locations restricted to customer points, or otherwise to a discrete set of points that is specified a priori?

We can see that our formulation has the least restrictive answers to these questions and thus addresses the problem in the greatest generality. That is, in our model we have the following.

- We allow for both the case that deliveries must be made by sidekicks and the case that the truck can also make deliveries.
- There can be any number of sidekicks on the truck and they are free to be launched and picked up in any order.
- The sidekick launch and pickup locations can be any point in the plane.

Other factors that distinguish the models surveyed here are the treatment of a restricted drone range and the way that the objective, be it completion time or cost or some measure of energy consumed, is determined. This work assumes unlimited drone range and that drones make a single delivery per trip from the truck. We take as our objective completion time. We assume that the time spent actually dropping a package at a customer node as well as the time spent capturing a sidekick and preparing it for relaunch are negligible. Both the truck and the sidekicks travel at fixed speeds along Euclidean distances. The specification of their relative speeds does however allow one to build some knowledge of the underlying network into the objective. One additional assumption that adds to the robustness of our formulation is that the sidekicks are allowed to be slower than the truck. We are thus able to accurately model systems like the truck-UGV schemes discussed in the introduction, whereas some papers surveyed require that the sidekicks be faster. For an in depth consideration of modeling concerns for drone routing problems see [20].

As this paper is concerned with theoretical analysis of sidekick problems, we deliberately narrow our focus, omitting a broader discussion of algorithmic approaches – whether exact or heuristic – in the interest of brevity. Wang et al. [30] consider the Vehicle Routing Problem with Drones in which multiple vehicles each carry multiple drones. They derive upper bounds on the improvement in completion time to be gained over the optimal TSP and VRP solutions without drones as well as the improvement to be gained by introducing faster drones. Poikonen et al. [21] extend the model of [30]. A battery life (time limit) is imposed on the drones; the possibility of using different distance metrics for the truck and drone and the possibility of using cost rather than time based objectives are considered; and there is an extension to the close-enough vehicle routing problem. Their results are bounds on improvement due to introduction of drones and due to different drone configurations.

Agatz et al. [2] produce a result that is a generalization of the results of [30] when applied to the TSP-D. That is, they give an upper bound on the improvement in completion time over just-truck routing allowing different distance metrics to be used for the truck and drone distances. The authors further give a lower bound to the TSP-D and an approximation algorithm using minimum spanning trees.

Campbell et al. [7] study a continuous approximation model for a sidekick problem with a truck carrying multiple drones. Demand is modeled as a continuous spatial density. Customer points are visited in rectangular swaths. The authors provide the expected cost of delivery in terms of the customer density and the truck and drone per-unit-distance and dropoff costs. Comparison is made to the expected cost without drones. Unlike in our model, drone launch and pickup locations are limited to customer points, and the sequence of deliveries is fixed to a truck delivery at which all drones are launched followed by another truck delivery at which all drones are picked up and relaunched.

Zhang [32] employs continuous approximation to characterize cost and emissions of truck-only, drone-only and truck-drone tandem delivery. Demand is modeled as a continuous spatial density. For the truck-drone tandem the truck carries a single drone which it launches at a truck-delivered customer point and retrieves and relaunches at its next truck delivery point. Thus the tour alternates truck and drone deliveries. The expected travel distances for a truck-drone route that visits customers in rectangular swaths is then determined and used to compute expected costs. Zhang's analysis and classification of multiple drone energy consumption models then informs modeling of drone emissions. This allows for computation of expected emissions for the different delivery systems. Zhang presents comparison, in cost and emissions, of the different delivery systems as well as comparison to the utilization of multiple systems for different subregions of the service region. Zhang shows how cost and emissions performance depend on drone and truck characteristics as well as the delivery density. Finally, an analysis of the tradeoff between cost and emissions is conducted.

In [8] Carlsson and Song consider the sidekick problem as formulated in this paper except restricted to only one sidekick and assuming that the sidekick is faster than the truck. Using a continuous approximation model that assumes a smooth demand distribution they are able to derive the asymptotic behavior of the optimal tour as the number of customers goes to infinity. This then yields a characterization of the improvement to be gained by introducing a sidekick and how this improvement depends on the relative speeds of the truck and sidekick.

3 Problem definitions

We begin by formally defining the problem of sidekick routing with multiple sidekicks. We assume that a single, uncapacitated truck must provide service to a collection of n customers in the plane, using the assistance of k sidekicks having unit capacity, and that the goal is to minimize the time to completion. To simplify exposition, we will first formulate our problem with an additional constraint that the truck itself is not permitted to visit any customers:

Definition 1. Let $p_1, \ldots, p_n$ be a collection of points in the plane. Let k denote the number of sidekicks. Let ϕ_0 denote the speed of the truck, and let ϕ_1 denote the speed of each sidekick (ϕ_1 can be greater or less than ϕ_0). Let variables $x_1, \ldots, x_n$ be the launch points for the sidekicks, and let variables $y_1, \ldots, y_n$ be the pickup points for the sidekicks. That is, point p_i is visited by a sidekick that is launched at point x_i and is retrieved at point y_i. Note that multiple launch/pickup events could occur at the same location, e.g. $x_i = y_j$.

Let variables z_j, $j \in \{1, \ldots, 2n\}$, be the location of the jth sidekick launch or pickup event, and let variables t_j, $j \in \{1, \ldots, 2n\}$, be the time of the jth sidekick launch or pickup event. The z_j's take the same values as the x_i's and the y_i's; we introduce them only to make indexing easier in the formulation. We let z_0 be the initial position of the truck and let z_{2n+1} be its final position. We require that the truck's tour be a loop, i.e. $z_{2n+1} = z_0$. We let t_0, equaling zero, be the time at which the truck starts its loop and let t_{2n+1} be the time at which the truck completes its loop.

Let $\sigma : \{1, \ldots, n\} \hookrightarrow \{1, \ldots, 2n\}$ map customer index i to the place of that customer's sidekick launch event in the ordering of all launch and pickup events. Similarly let $\pi : \{1, \ldots, n\} \hookrightarrow \{1, \ldots, 2n\}$ map customer index i to the place of that customer's sidekick pickup event in the ordering of all launch and pickup events. That is if $\sigma(3) = 5$ and $\pi(3) = 8$, customer 3 is serviced by a sidekick whose launch is the 5th launch/pickup event and whose pickup is the 8th launch/pickup event. Let $\mathcal{F}$ be the set of all pairs of mappings (σ, π) that induce a valid sidekick tour. The conditions for inclusion in $\mathcal{F}$ are

$$\sigma(i) < \pi(i) \quad \forall i \in \{1, \ldots, n\} \qquad \text{(launches occur before corresponding pickups)}$$
$$|\{i : \sigma(i) < j\}| - |\{i : \pi(i) < j\}| \le k \quad \forall j \qquad \text{(never more than } k \text{ sidekicks in use)}$$
$$\sigma, \pi \text{ are injective} \qquad \text{(a launch and a pickup for each customer)}$$
$$\sigma(i) \ne \pi(i') \quad \forall i, i' \in \{1, \ldots, n\}. \qquad \text{(one event per place in the ordering)}$$

The second condition says that at any time there have been at most k more launch events than there have been pickup events, ensuring that we are never making use of more than k sidekicks at a time. The last two conditions say that the maps σ and π have to jointly form a bijection between $\{1, \ldots, n\}$ and $\{1, \ldots, 2n\}$ (to be precise, the two maps actually form a bijection between the multiset $\{1, \ldots, n\} \uplus \{1, \ldots, n\}$ and $\{1, \ldots, 2n\}$, where $\uplus$ denotes the multiset union [14]).

The *sidekick problem* problem is then given by

$$\underset{\boldsymbol{x},\boldsymbol{y},\boldsymbol{z},\boldsymbol{t},\sigma,\pi}{\text{minimize}} \quad t_{2n+1} \quad s.t. \tag{SK1}$$
$$t_j \ge t_{j-1} + \frac{1}{\phi_0}\|z_j - z_{j-1}\| \quad \forall j \in \{1, \ldots, 2n+1\} \tag{1}$$
$$t_{\pi(i)} \ge t_{\sigma(i)} + \frac{1}{\phi_1}\|x_i - p_i\| + \frac{1}{\phi_1}\|p_i - y_i\| \quad \forall i \in \{1, \ldots, n\} \tag{2}$$
$$z_{\sigma(i)} = x_i \quad \forall i \in \{1, \ldots, n\}$$
$$z_{\pi(i)} = y_i \quad \forall i \in \{1, \ldots, n\}$$
$$t_0 = 0$$
$$z_{2n+1} = z_0$$
$$(\sigma, \pi) \in \mathcal{F},$$

where the objective value is the time at which the truck completes its loop, (1) captures the time needed for the truck to travel between launch and pickup points, and (2) captures the time needed for a sidekick to travel from its launch point, to a customer, and then to its pickup point.

To extend (SK1) to the case where the truck is permitted to visit customers, some additional notation is required:

Definition 2. We partition the set of customers into two sets $\mathcal{S} \subseteq \{1, \ldots, n\}$, representing those customers visited by a sidekick, and its complement $\mathcal{T} = \bar{\mathcal{S}}$, representing those customers visited by the truck (these sets are optimization variables because we can choose which customers are visited by the truck). The number of events is now equal to $m := 2|\mathcal{S}| + |\mathcal{T}|$ because a truck visiting a customer counts as only one event. This necessitates a third map $\theta : \mathcal{T} \hookrightarrow \{1, \ldots, m\}$, in addition to the maps $\sigma, \pi : \mathcal{S} \hookrightarrow \{1, \ldots, m\}$. Let $\mathcal{F}$ be the set of all (σ, π, θ) that induce a valid sidekick tour. We have the same conditions as in the previous problem that ensure σ and π do not pickup before launching or use more than k sidekicks. In addition, in this case we must require that each sidekick-visited customer has a launch and a pickup event and each truck-visited customer has a truck visit event, with each of

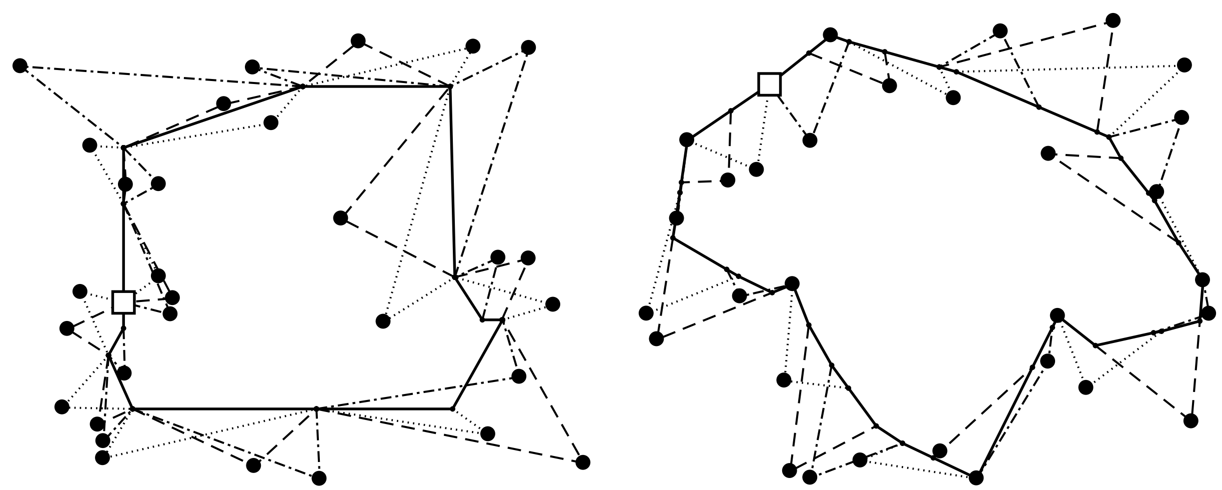

(a)

(b)

Figure 2: Solutions to the sidekick problem for $n = 30$ customers, $k = 3$ sidekicks, and a sidekick speed, ϕ_1, that is twice the speed, ϕ_0, of the truck. The solid line is the truck tour; each sidekick's route has its own dashed/dotted line style; the square represents the starting and finishing point of the truck. In the figure on the left all deliveries are made by the sidekicks (Problem SK1). In the figure on the right the truck is also allowed to make deliveries (Problem SK2).

these events being mapped to a unique place in the ordering of events. That is,

$$\sigma, \pi : \mathcal{S} \hookrightarrow \{1, \ldots, m\}$$
$$\theta : \mathcal{T} \hookrightarrow \{1, \ldots, m\}$$
$$\sigma(\mathcal{S}), \pi(\mathcal{S}), \theta(\mathcal{T}) \text{ are pairwise disjoint.}$$

Put another way, the maps σ, π, and θ have to jointly form a bijection between $\mathcal{S} \cup \mathcal{T}$ and $\{1, \ldots, m\}$ (to be precise, the three maps actually form a bijection between the multiset $\mathcal{S} \uplus \mathcal{S} \cup \mathcal{T}$ and $\{1, \ldots, m\}$).

The extension to (SK1) is then a natural one:

$$\underset{\boldsymbol{x},\boldsymbol{y},\boldsymbol{z},\boldsymbol{t},\sigma,\pi,\theta}{\text{minimize}} \quad t_{m+1} \quad s.t. \tag{SK2}$$

$$t_{j+1} \geq t_j + \frac{1}{\phi_0}\|z_{j+1} - z_j\| \quad \forall j \in \{1, \dots, m\} \tag{3}$$

$$t_{\pi(i)} \geq t_{\sigma(i)} + \frac{1}{\phi_1}(\|x_i - p_i\| + \|p_i - y_i\|) \quad \forall i \in \mathcal{S} \tag{4}$$

$$z_{\sigma(i)} = x_i \quad \forall i \in \mathcal{S}$$

$$z_{\pi(i)} = y_i \quad \forall i \in \mathcal{S}$$

$$z_{\theta(i)} = p_i \quad \forall i \in \mathcal{T}$$

$$t_0 = 0$$

$$z_{m+1} = z_0$$

$$(\sigma, \pi, \theta) \in \mathcal{F},$$

where $\mathcal{S}$ is defined as the domain of variables σ and π and $\mathcal{T}$ is the domain of variable θ.

Figure 2 shows examples of solutions to the problems defined above for 30 customers with multiple sidekicks that are faster than the truck.

4 Preliminaries

Having defined two variants of sidekick routing, we now turn to some preliminary results that will be useful in our analysis of these problems. This section presents existing results from prior work as well as some additional analysis of our own.

4.1 Existing results from related work

The concept of a *subadditive Euclidean functional* was introduced in [25], which provides a key insight that we will use in this paper:

Definition 3. A function $L(\cdot)$ from the set of finite subsets of $\mathbb{R}^2$ to the non-negative real numbers is said to be a *monotone subadditive Euclidean functional on* $\mathbb{R}^2$ if it satisfies the following properties:

1. $L(\emptyset) = 0$.
2. *Homogeneity*: $L(\alpha x_1, \dots, \alpha x_n) = \alpha L(x_1, \dots, x_n)$ for all real $\alpha > 0$.
3. *Translation invariance*: $L(x_1 + x, \dots, x_n + x) = L(x_1, \dots, x_n)$ for all $x \in \mathbb{R}^2$.
4. *Monotonicity*: $L(x \cup A) \geq L(A)$ for any $x \in \mathbb{R}^2$ and finite subset $A \subset \mathbb{R}^2$.
5. *Geometric subadditivity*: There exists a constant $C > 0$, such that for all positive integers m, n

and $\{x_1, \ldots, x_n\} \in [0,1]^2$, we have

$$L(x_1, \ldots, x_n) \leq \sum_{i=1}^{m^2} L(\{x_1, \ldots, x_n\} \cap Q_i) + Cm$$

where $\{Q_i\}$, $1 \leq i \leq m^2$ is the partition of $[0,1]^2$ into squares of edge length $1/m$.

Examples of subadditive Euclidean functionals include the TSP tour and the Steiner tree. The minimum spanning tree, the minimum matching, and the nearest neighbor graph are all "close" to being subadditive Euclidean functionals, but violate the monotonicity requirement (though it turns out that this can easily be overcome for all relevant applications). The monographs [26, 31] are devoted to more general settings for Theorem 5, with the most prominent generalization being the following:

Theorem 4 (basic theorem of subadditive Euclidean functionals)**.** *Suppose L is a monotone subadditive Euclidean functional defined on $\mathbb{R}^2$. If the random variables $\{X_i\}$ are independent with the uniform distribution on $[0,1]^2$, then with probability one, we have*

$$\frac{L(X_1, \ldots, X_n)}{\sqrt{n}} \to \beta_L$$

as $n \to \infty$, where $\beta_L \geq 0$ is a constant.

The above is a generalization of the following classical theorem, originally stated in [5] and further developed in [25, 26], is one of the fundamental results of the continuous approximation paradigm; it relates the length of a TSP tour of a sequence of points to the distribution from which they were sampled:

Theorem 5 (BHH Theorem)**.** *Suppose that $X_1, X_2, \ldots$ is a sequence of random points i.i.d. according to an absolutely continuous probability density function f defined on a compact planar region $\mathcal{R}$. Then with probability one, the length $\mathsf{TSP}(X_1, \ldots, X_n)$ of the optimal travelling salesman tour through all X_i's satisfies*

$$\lim_{n \to \infty} \frac{\mathsf{TSP}(X_1, \ldots, X_n)}{\sqrt{n}} = \beta_{\text{TSP}} \iint_{\mathcal{R}} \sqrt{f(x)}\, \mathrm{d}x$$

where β_{TSP} is a positive constant.

Although the exact value of β_{TSP} is unknown, it has been shown that $0.6277 \leq \beta_{\text{TSP}} \leq 0.9204$; see [3, 5, 11].

We conclude with some additional problem definitions and convergence results that will also prove key to our analysis:

Definition 6 (Medians Problem)**.** Given a collection of points $x_1, \ldots, x_n$ in $\mathbb{R}^2$ and a positive integer p, the the *p-medians problem* is given by

$$\mathsf{PMed}(x_1, \ldots, x_n; p) := \min_{\mathcal{S} \subset \{1, \ldots, n\} : |S| \leq p} \sum_{i=1}^{n} \min_{j \in \mathcal{S}} \|x_i - x_j\|;$$

that is, the problem of selecting a subset $\mathcal{S} \subset \{1, \dots, n\}$ of *median* points such that $|\mathcal{S}| \leq p$, that minimizes the sum of the distances from all points to their nearest median.

Definition 7 (Balanced Medians Problem)**.** The *balanced medians problem* $\mathsf{BMed}(x_1, \dots, x_n; d)$ is a further-constrained variation of the p-medians problem. We can equivalently express p-medians as the problem of selecting a set of medians $\mathcal{S} \subset \{1, \dots, n\}$ and an assignment of the points x_i to medians such that the sum of the distances from the points to their assigned medians is minimized. With no constraint on our assignment selection we have that in the p-medians problem the optimal assignment for any median set is simply to assign a point to its nearest median. The balanced medians problem imposes an additional constraint on the assignment selection, namely median $x_j \in \mathcal{S}$ can have at most $d \geq 2$ non-median points assigned to it. It is further required that each median is assigned to itself.

That is,

$$\mathsf{BMed}(x_1, \dots, x_n; d) := \min_{\substack{\mathcal{S} \subset \{1,\dots,n\}\colon |\mathcal{S}| = p \\ \mu\colon \{1,\dots,n\} \mapsto \mathcal{S}}} \sum_{i=1}^{n} \|x_i - x_{\mu(i)}\|, \tag{5}$$

where

$$p = \left\lceil \frac{n}{d+1} \right\rceil,$$

$x_{\mu(i)}$ is the median assigned to point x_i, and for all j such that $j \in \mathcal{S}$, $x_{\mu(j)} = x_j$ and $x_{\mu(i)} = x_j$ for at most d of the $i \neq j$.

The following result is due to [16]:

Theorem 8 (Asymptotic convergence of the balanced medians problem)**.** *The balanced medians problem satisfies the same convergence as in Theorem 4; that is, if $X_1, X_2, \dots$ is a sequence of random points i.i.d. according to an absolutely continuous probability density function f defined on a compact planar region $\mathcal{R}$ and $d \geq 2$ is fixed, then with probability one, the cost $\mathsf{BMed}(X_1, \dots, X_n; d)$ satisfies*

$$\lim_{n \to \infty} \frac{\mathsf{BMed}(X_1, \dots, X_n; d)}{\sqrt{n}} = \beta_{\mathsf{BMed}}(d) \iint_{\mathcal{R}} \sqrt{f(x)}\, \mathrm{d}x$$

where $\beta_{\mathsf{BMed}}(d)$ depends only on d.

4.2 Further notes on Theorem 8

This section describes a lower bound on the function $\beta_{\mathsf{BMed}}(d)$ from Theorem 8.

Theorem 9. *The function $\beta_{\mathsf{BMed}}(d)$ satisfies*

$$\beta_{\mathsf{BMed}}(d) \geq \frac{\sqrt{2} d (d!)^{1/(2d)}}{\sqrt{\pi e} (d+1)^{(2d+1)/(2d)}}.$$

That is, with probability one,

$$\lim_{n \to \infty} \frac{\mathsf{BMed}(X_1, \dots, X_n; d)}{\sqrt{n}} \geq \frac{\sqrt{2} d (d!)^{1/(2d)}}{\sqrt{\pi e} (d+1)^{(2d+1)/(2d)}} \iint_{\mathcal{R}} \sqrt{f(x)}\, \mathrm{d}x.$$

Proof. See Section A of the Online supplement. □

Remark 10. Using Stirling's approximation to simplify the factorial and taking logarithms, it is routine to verify that the lower bound above satisfies

$$\frac{\sqrt{2}d(d!)^{1/(2d)}}{\sqrt{\pi e}(d+1)^{(2d+1)/(2d)}} \geq 0.2886\sqrt{d}$$

for $d \geq 100$ (the choice of 100 is merely an arbitrary "large number").

5 A continuous approximation analysis

This section describes a continuous approximation analysis of the sidekick routing problems (SK1) and (SK2).

5.1 Naive asymptotic analysis

Relying solely on Theorem 4, we can obtain the following partial characterization of the asymptotic behavior of both problems (SK1) and (SK2).

Claim 11. For fixed values of k, ϕ_0, and ϕ_1, let $T(p_1, \ldots, p_n)$ denote the optimal objective value of problem (SK1). Then if the customer points p_i consist of random samples P_i independently drawn from a uniform distribution on the unit square, then there exists a non-negative constant $c_{\text{SK1}} = c_{\text{SK1}}(k, \phi_0, \phi_1)$ such that

$$\frac{T(P_1, \ldots, P_n)}{\sqrt{n}} \to c_{\text{SK1}}$$

with probability one as $n \to \infty$. The same statement holds when $T(\cdot)$ is the optimal objective value of problem (SK2), with a different constant $c_{\text{SK2}} \leq c_{\text{SK1}}$.

Proof. This follows immediately from Theorem 4 because $T(\cdot)$ is a monotone subadditive Euclidean functional as defined in Definition 3. We verify that $T(\cdot)$ meets the definition.

1. Clearly $T(\emptyset) = 0$.

2. Given $\alpha > 0$ scaling all customer point locations by α scales all travel times by α and thus the optimal solution to problems $SK1$ and $SK2$ remains the same up to scaling and the total travel time for the solution scales by α.

3. It is clear that translating all customer locations by the same x will not impact the solution to either sidekick problem.

4. Adding a customer that must be visited cannot reduce the total time of the route.

5. Finally, suppose we have solutions to the sidekick problem for the customer points within each square Q_i in $\{Q_i\}$, $1 \leq i \leq m^2$, a partition of $[0,1]^2$ into squares of edge length $1/m$. Then we can stitch together these solutions to obtain a feasible solution on all points $\{x_1, \ldots, x_n\}$ in

$[0,1]^2$. For each square we simply perform the in-square route, then have the truck drive from the truck starting and ending location for the square to the truck starting and ending location for the adjacent square. Note that the truck has all the sidekicks on it at the starting and ending locations, so this yields a valid tour. The objective value (i.e. the makespan) of our solution will be the sum of the in-square travel times $T(\{x_1,\ldots,x_n\}\cap Q_i)$ plus the between-square driving time, which is at most the number of squares times $2\sqrt{2}/(m\phi_0)$. That is, the objective value is less than or equal to $T(\{x_1,\ldots,x_n\}\cap Q_i)+2\sqrt{2}m/\phi_0$. The optimal solution on $\{x_1,\ldots,x_n\}$ must have objective value less than this feasible solution.

□

Claim 11 describes the scaling behavior of our problem as $n\to\infty$, namely that the objective value scales proportionally to $\sqrt{n}$, but it tells us nothing about c_{SK1} (or c_{SK2}). For example, it is obvious that both are non-increasing with respect to the three fixed parameters ϕ_0, ϕ_1, and k (since making things faster or increasing the number of sidekicks cannot possibly make the process slower), and routine scaling arguments establish that $c_{\mathrm{SK1}}(k,\phi_0,\phi_1)=\phi_0 c_{\mathrm{SK1}}(k,1,\phi_1/\phi_0)$ for all k,ϕ_0,ϕ_1 (and similarly for c_{SK2}). We devote the remainder of this section to a more precise analysis of c_{SK1} and c_{SK2}.

5.2 A lower bound for (SK2)

Of course, problem (SK2) is itself a lower bound of (SK1) by construction, so it will suffice to consider (SK2) only. This bound is of course poor when $\phi_1\ll\phi_0$, because slow sidekicks will result in more work imposed on the host vehicle, ceteris paribus. We derive a lower bound for (SK2) in terms of the Traveling Salesman tour and solution to a balanced medians problem on the p_i.

Lemma 12. *Let T_n denote the optimal objective value for Problem (SK2). We have*

1. $\mathsf{TSP}(p_1,\ldots,p_n)\le(\phi_0+k\phi_1)T_n$
2. $\mathsf{BMed}(p_1,\ldots,p_n;d)\le(d\phi_0+k\phi_1)T_n$ for all $d\ge 2$.

Proof. For the first claim, we can construct a TSP solution from the (SK2) solution as follows. Consider an optimal solution $(\mathbf{x},\mathbf{y},\mathbf{z},\mathbf{t},\sigma,\pi,\theta,\mathcal{S},\mathcal{T})$ to (SK2). For each $i\in\mathcal{S}$, the set of customers visited by the sidekicks, let

$$u_i := \operatorname*{argmin}_{u\in\{x_i,y_i\}} ||u-p_i||,$$

that is u_i is the closer to the customer of its sidekick launch and pickup points. For each $i\in\mathcal{T}$ let

$$u_i := p_i.$$

We then construct a TSP tour of the points as follows. Let the tour follow the path of the truck, visiting the customers in $\mathcal{T}$ along the tour. Whenever we reach one of the u_i for $i\in\mathcal{S}$, let the tour travel from u_i to p_i and back, then continue along the truck path. It is clear that the length added to our TSP tour coming from the truck's route is less than or equal to $\phi_0 T_n$, the truck's speed times the total time for our sidekick tour.

To bound the length from visiting points in $\mathcal{S}$ we let P_j be the set of points visited by sidekick j. Then

$$T_n \geq \frac{1}{\phi_1} \sum_{i:p_i \in P_j} ||x_i - p_i|| + ||p_i - y_i|| \quad \forall j \in \{1, \ldots, k\}.$$

That is the total sidekick tour time exceeds the time any given sidekick travels. Dividing the above by k and summing over all j yields

$$\begin{aligned} T_n &\geq \frac{1}{k\phi_1} \sum_{i \in \mathcal{S}} ||x_i - p_i|| + ||p_i - y_i|| \\ &\geq \frac{2}{k\phi_1} \sum_{i \in \mathcal{S}} ||u_i - p_i||. \end{aligned}$$

Twice the sum of the $||u_i - p_i||$ is precisely what we add to our TSP tour to visit $\mathcal{S}$. Thus the contribution of this part of our TSP tour is bounded by $k\phi_1 T_n$. Adding together our truck and sidekick pieces of the TSP tour and applying the triangle inequality gives the result.

For the second bound we can construct a balanced median solution from the (SK2) solution as follows. Think of the truck as completing a tour on our u_i defined as above. Group every $d+1$ of the customer points associated with the u_i along this tour and choose as their median the point which is closest to the tour. This construction is pictured in Figure 4. By the triangle inequality, the distance from a point to its assigned median is less than or equal to the distance of traveling from that point to its corresponding u_i, then traveling along the truck tour to the median's corresponding u_i, then traveling out to the assigned median. The cost of this balanced medians solution, i.e. sum of these distances, is then less than or equal to the sum of the distances from the non-median points to their corresponding u_i, plus d times the length of the tour of the u_i, plus the sum, over all medians, of d times the distance from the median's corresponding u to the median. By our selection of the medians it is clear that this last sum is less than or equal to the sum of all of the distances from non-median points to their corresponding u. Then, noting once again that

$$\phi_0 T_n \geq \text{truck tour of the } u_i,$$

and

$$\begin{aligned} \frac{k\phi_1}{2} \cdot T_n &\geq \sum_{i \in \mathcal{S}} ||u_i - p_i|| \\ &= \sum_{i=1}^{n} ||u_i - p_i||, \qquad (u_i = p_i \text{ for } i \in \mathcal{T}) \end{aligned}$$

the length of the balanced medians solution is less than or equal to

$$\left(\frac{k\phi_1}{2} + d\phi_0 + \frac{k\phi_1}{2} \right) T_n.$$

□

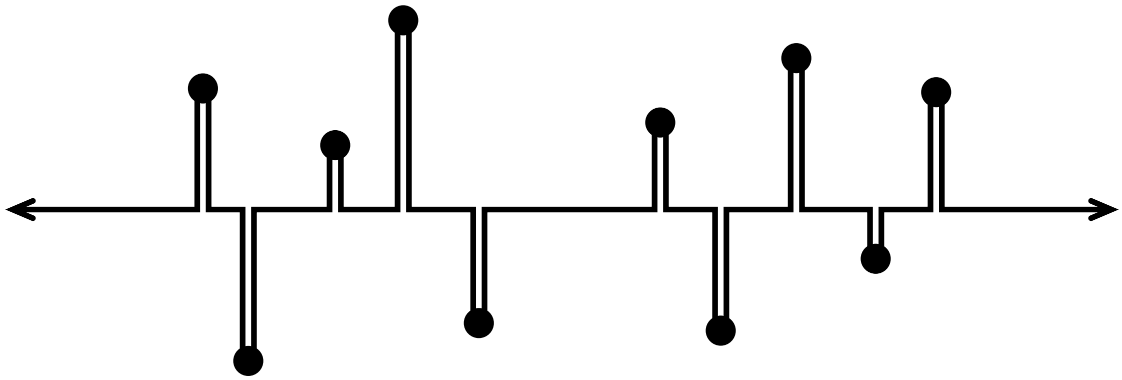

Figure 3: Constructing a TSP solution from a solution to problem (SK2). The horizontal line represents the tour of the u_i (the closer of the launch and pickup points for sidekick-visited customers and the p_i for truck-visited customers) in the (SK2) solution. We follow the tour, traveling from u_i to p_i and back for each $i \in \mathcal{S}$. If T_n is the objective value of the problem (SK2) solution then the total cost of the resulting TSP solution is less than or equal to $(\phi_0 + k\phi_1)T_n$.

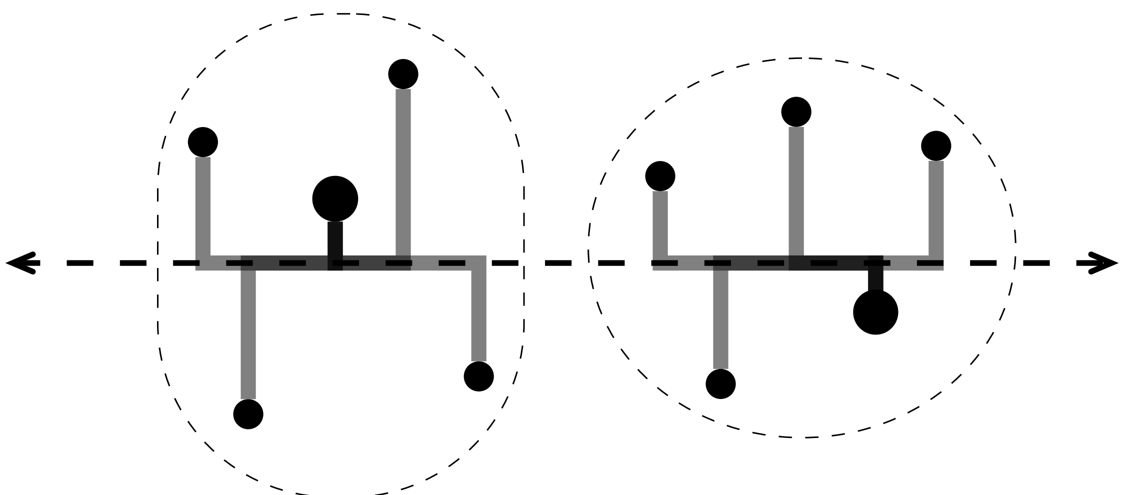

Figure 4: Constructing a balanced medians solution from a solution to problem (SK2). The horizontal line represents the tour of the u_i (the closer of the launch and pickup points for sidekick-visited customers and the p_i for truck-visited customers) in the (SK2) solution. Here we choose $d = 4$ and group every 5 points along the tour. We choose as the median for these 5 points the point which is closest to the tour. Using the paths pictured, it is clear that to connect all points to their medians we need travel at most d times the length of the truck tour plus twice the total distance from the points to their u_i. If T_n is the objective value of the problem (SK2) solution then the total cost of the resulting Bounded Medians solution is less than or equal to $(d\phi_0 + k\phi_1)\, T_n$.

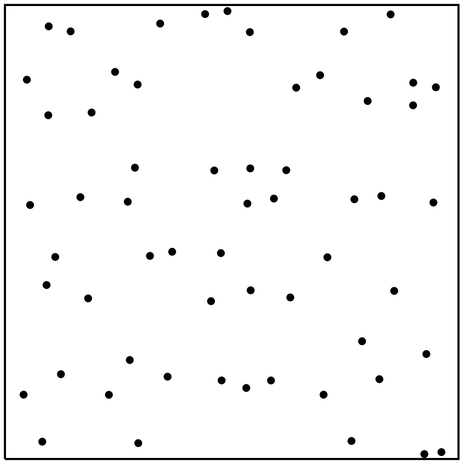
(a) Input

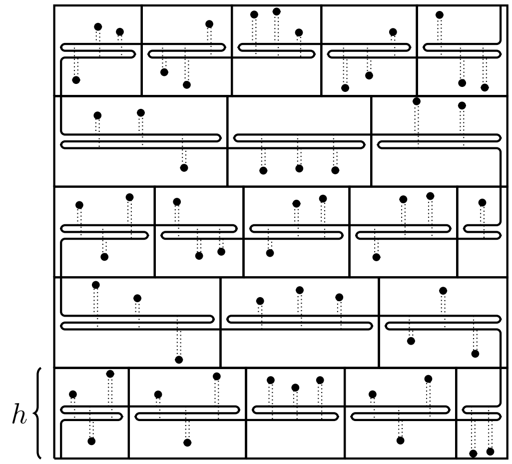

(b) Output

Figure 5: The tour described in Lemma 13, assuming $k = 3$.

5.3 An upper bound for (SK1)

To bound the objective value of (SK1), we describe a simple "zig-zagging" heuristic in the unit square:

Lemma 13. *For fixed ϕ_0, ϕ_1, and k and points $p_1, \dots, p_n$ lying in the unit square, there exists a routing strategy for problem (SK1) whose time to completion $T(p_1, \dots, p_n)$ satisfies*

$$T(p_1, \dots, p_n) \leq \frac{2\sqrt{3}}{\sqrt{\phi_0 \phi_1 k}} \cdot \sqrt{n} + C$$

where C is a constant that depends only on ϕ_0, ϕ_1, and k.

Proof. Divide the unit square into strips of height $h = \sqrt{3\phi_1 k/(\phi_0 n)}$ (there may be one strip whose height is less than this due to rounding). There are $m = \left\lceil \sqrt{\phi_0 n/(3\phi_1 k)} \right\rceil \leq \sqrt{\phi_0 n/(3\phi_1 k)} + 1$ such strips. Further subdivide each strip into rectangles so that each rectangle (except possibly the right-most in each strip) contains k points. There are at most $m + n/k$ rectangles in total. Finally, construct a tour for the truck and all sidekicks by traversing each rectangle three times, releasing the sidekicks on the first traversal and retrieving the sidekicks on the third traversal, as illustrated in Figure 5.

It is easy to see that for a rectangle having width w (and height h), it is possible to perform three horizontal traversals and release and retrieve the sidekicks in at most $3w/\phi_0 + h/\phi_1$ time units. We release the sidekicks when vertically aligned with a customer on the first left-to-right traversal, make a right-to-left traversal and then, when the sidekicks have all returned to the middle line, make another left-to-right traversal to retrieve them and continue on to the next rectangle. It is also easy to see that the only remaining time needed is for the truck to perform vertical moves to move from one strip to the next, which is a constant amount of $1/\phi_0$ time units, plus whatever time is needed for the truck to return to its point of origin, which is also at most $\sqrt{2}/\phi_0$ time units. Hence, if we let w_i denote the

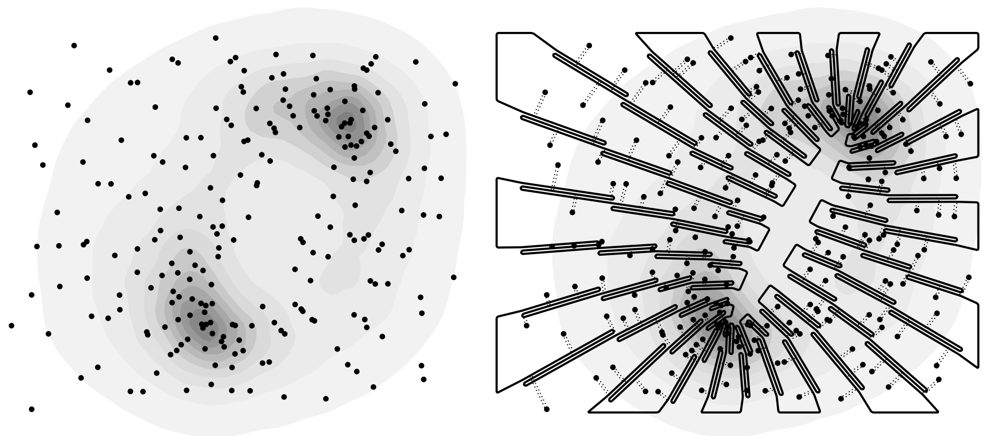

(a) yo (b) yo

Figure 6: The tour described in Lemma 14; it consists of the same kind of tour as in Figure 5, but "scaled" with respect to the probability distribution that is indicated by shading.

width of rectangle i, then the total amount of time to complete this tour is at most

$$\begin{aligned}(1+\sqrt{2})/\phi_0 + \sum_i (3w_i/\phi_0 + h/\phi_1) &\leq (1+\sqrt{2})/\phi_0 + \frac{3}{\phi_0}\underbrace{\sum_i w_i}_{=m} + (m+n/k)h/\phi_1 \\ &\leq (1+\sqrt{2})/\phi_0 + \frac{3}{\phi_0}\left(\sqrt{\frac{\phi_0 n}{3\phi_1 k}}+1\right) + \frac{1}{\phi_1}\left(\sqrt{\frac{\phi_0 n}{3\phi_1 k}}+1+n/k\right)\sqrt{\frac{3\phi_1 k}{\phi_0 n}} \\ &= \frac{2\sqrt{3}}{\sqrt{\phi_0\phi_1 k}}\cdot\sqrt{n} + \sqrt{\frac{3k}{\phi_0\phi_1 n}} + \frac{1}{\phi_1} + (4+\sqrt{2})/\phi_0\end{aligned}$$

as desired. □

Lemma 13 is deterministic, but also implies the following:

Lemma 14. *Let ϕ_0, ϕ_1, and k be fixed and let $P_1, \ldots, P_n$ be independent samples from an absolutely continuous probability density f with compact support $\mathcal{R}$. The optimal time to completion $T(P_1, \ldots, P_n)$ for problem (SK1) satisfies*

$$\limsup_{n\to\infty} \frac{T(P_1,\ldots,P_n)}{\sqrt{n}} \leq \frac{3.47}{\sqrt{\phi_0\phi_1 k}} \cdot \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x$$

with probability one.

Proof. This is a routine scaling argument, together with the law of large numbers and the fact $T(\cdot)$ is a subadditive Euclidean functional (see Claim 11); see Section B of the Online Supplement for details. □

Remark 15. Another routing strategy is to subdivide the rectangles as before, but to release all sidekicks simultaneously at one end of the rectangle, perform only one horizontal traversal, and rendezvous

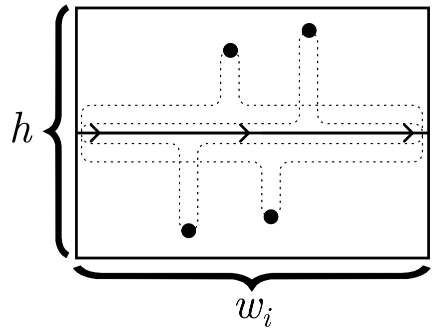

Figure 7: Using the alternate routing strategy from Remark 15, the truck traverses each rectangle only one time (in this case from left to right as indicated by the arrows), releasing the k sidekicks at one end of the rectangle and retrieving them at the other.

with the sidekicks at the other end of the rectangle, as shown in Figure 7. It is obvious that the amount of time to cover each rectangle is at most

$$\max\left\{\frac{w_i}{\phi_0}, \frac{w_i+h}{\phi_1}\right\} \le \frac{w_i}{\phi_0} + \frac{w_i+h}{\phi_1}\,,$$

and it is routine to verify that under this policy, we have

$$\limsup_{n\to\infty} \frac{T(P_1,\ldots,P_n)}{\sqrt{n}} \le 2\sqrt{\frac{\phi_0+\phi_1}{\phi_0\phi_1^2 k}} \cdot \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x$$

as opposed to the coefficient $3.47/\sqrt{\phi_0\phi_1 k}$ of Lemma 14. This policy is preferable whenever $\phi_0 < 2\phi_1$ (without any assumptions on k), and approaches $2/\sqrt{\phi_0\phi_1 k}$ (i.e. a 57% reduction from our original bound) when $\phi_1 \gg \phi_0$.

5.4 Convergence analysis for (SK1) and (SK2)

We have now collected enough supporting evidence for our main claim:

Theorem 16. *Let ϕ_0, ϕ_1, and k be fixed. Let T_n denote the optimal objective value to problem (SK1), where input points $P_1,\ldots,P_n$ are independent uniform samples in the unit square. Then there exists a constant β_{SK1} satisfying $0.1368 < \beta_{\mathrm{SK1}} < 3.47$ such that*

$$\frac{T_n}{\sqrt{n}} \to \frac{\beta_{\mathrm{SK1}}}{\sqrt{\phi_0 \max\{\phi_0, \phi_1 k\}}} \tag{6}$$

with probability one as $n\to\infty$. Moreover, the same statement holds for a different constant $\beta_{\mathrm{SK2}} \le \beta_{\mathrm{SK1}}$ when T_n is the optimal objective value to problem (SK2), which also satisfies $0.1368 < \beta_{\mathrm{SK2}} < 3.47$. Finally, when the points $P_1,\ldots,P_n$ are independent samples from an absolutely continuous probability density f with compact support $\mathcal{R}$, we have

$$0.1368c \le \liminf \frac{T_n}{\sqrt{n}} \le \limsup \frac{T_n}{\sqrt{n}} \le 3.47c$$

with probability one as $n \to \infty$, where T_n is the optimal objective value to either problem (SK1) or (SK2), and

$$c = \frac{\iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x}{\sqrt{\phi_0 \max\{\phi_0, \phi_1 k\}}}\,.$$

Proof. To simplify notation, we introduce the parameter t defined as

$$t = \phi_1 k/\phi_0$$

throughout this proof, and rewrite the desired result (6) equivalently as

$$\phi_0 \sqrt{\max\{1,t\}} \cdot \frac{T_n}{\sqrt{n}} \to \beta_{\mathrm{SK1}}\,.$$

The existence of β_{SK1} and β_{SK2} was already established in Claim 11 (set $\beta_{\mathrm{SK1}} = c_{\mathrm{SK1}}\sqrt{\phi_0 \max\{\phi_0, \phi_1 k\}}$ and so forth); the real work lies in computing the bounds on these constants. Since $\beta_{\mathrm{SK2}} \le \beta_{\mathrm{SK1}}$, it will suffice to show that $0.1368 < \beta_{\mathrm{SK2}}$ and that $\beta_{\mathrm{SK1}} < 3.47$. To show that $0.1368 < \beta_{\mathrm{SK2}}$, Lemma 12 says that

$$T_n \ge \frac{\mathsf{TSP}(P_1,\dots,P_n)}{\phi_0 + t\phi_0} \tag{7}$$

$$\implies \lim_{n\to\infty} \frac{T_n}{\sqrt{n}} \ge \lim_{n\to\infty} \frac{\mathsf{TSP}(P_1,\dots,P_n)}{\phi_0(1+t)\sqrt{n}} = \frac{\beta_{\mathrm{TSP}}}{\phi_0(1+t)} \tag{8}$$

$$\implies \phi_0\sqrt{\max\{1,t\}} \cdot \lim_{n\to\infty} \frac{T_n}{\sqrt{n}} \ge \phi_0\sqrt{\max\{1,t\}} \cdot \frac{\beta_{\mathrm{TSP}}}{\phi_0(1+t)} \ge \frac{0.6277\sqrt{\max\{1,t\}}}{1+t}$$

$$\implies \beta_{\mathrm{SK2}} \ge \frac{0.6277\sqrt{\max\{1,t\}}}{1+t} > 0.1368 \text{ whenever } t < 19, \tag{9}$$

where in (8) we are justified in taking limits as we have seen such limits exist for problem (SK2) (Claim 11) and for the TSP (Theorem 5).

In addition Lemma 12 tells us that, provided $t \ge 2$,

$$T_n \ge \frac{\mathsf{BMed}(P_1,\dots P_n; \lfloor t \rfloor)}{\lfloor t \rfloor \phi_0 + t\phi_0} \tag{10}$$

$$\implies \lim_{n\to\infty} \frac{T_n}{\sqrt{n}} \ge \lim_{n\to\infty} \frac{\mathsf{BMed}(P_1,\dots P_n; \lfloor t \rfloor)}{\phi_0(\lfloor t \rfloor + t)\sqrt{n}} = \frac{\beta_{\mathsf{BMed}}(\lfloor t \rfloor)}{\lfloor t \rfloor + t} \tag{11}$$

$$\implies \phi_0\sqrt{\max\{1,t\}} \cdot \lim_{n\to\infty} \frac{T_n}{\sqrt{n}} \ge \frac{\phi_0 \beta_{\mathsf{BMed}}(\lfloor t \rfloor)\sqrt{\max\{1,t\}}}{\phi_0(\lfloor t \rfloor + t)} \ge \frac{\sqrt{2}\lfloor t \rfloor (\lfloor t \rfloor !)^{1/(2\lfloor t \rfloor)}\sqrt{\max\{1,t\}}}{\sqrt{\pi e}(\lfloor t \rfloor + 1)^{(2\lfloor t \rfloor+1)/(2\lfloor t \rfloor)}(\lfloor t \rfloor + t)} \tag{12}$$

$$\implies \beta_{\mathrm{SK2}} \ge \frac{\sqrt{2}\lfloor t \rfloor (\lfloor t \rfloor !)^{1/(2\lfloor t \rfloor)}\sqrt{\max\{1,t\}}}{\sqrt{\pi e}(\lfloor t \rfloor + 1)^{(2\lfloor t \rfloor+1)/(2\lfloor t \rfloor)}(\lfloor t \rfloor + t)} > 0.1368 \text{ whenever } t \ge 19, \tag{13}$$

where in (11) we are justified in taking limits as we have seen such limits exists for problem (SK2) (Claim 11) and for the balanced medians problem (Theorem 8), and in (12) we have applied Theorem 9.

The upper bound $\beta_{\text{SK1}} < 3.47$ is very simple. From Lemma 13, we have

$$\begin{aligned} T_n &\le \frac{2\sqrt{3}}{\phi_0\sqrt{t}} \cdot \sqrt{n} + C \qquad (14) \\ \implies \phi_0\sqrt{\max\{1,t\}} \cdot \lim_{n\to\infty} \frac{T_n}{\sqrt{n}} &\le \phi_0\sqrt{\max\{1,t\}} \cdot \lim_{n\to\infty}\left(\frac{2\sqrt{3}}{\phi_0\sqrt{t}} + \frac{C}{\sqrt{n}}\right) \\ \implies \beta_{\text{SK1}} &\le 2\sqrt{3} < 3.47 \text{ for } t \ge 1, \end{aligned}$$

and for $t < 1$, we simply eschew the sidekicks altogether and visit all of the P_i's with the truck (to be precise, since the truck is not allowed to visit any points in (SK1), we bring the truck within arbitrarily small distance ϵ from each P_i and release and retrieve one of the sidekicks):

$$\begin{aligned} \lim_{n\to\infty} \frac{\phi_0 T_n}{\sqrt{n}} &\le \beta_{\text{TSP}} \\ \implies \phi_0\sqrt{\max\{1,t\}} \cdot \lim_{n\to\infty} \frac{T_n}{\sqrt{n}} &\le \phi_0\sqrt{\max\{1,t\}} \cdot \frac{\beta_{\text{TSP}}}{\phi_0} = \beta_{\text{TSP}} \\ \implies \beta_{\text{SK1}} &\le \beta_{\text{TSP}} \le 0.9204 \text{ for } t < 1 \end{aligned}$$

as desired. This completes the proof of the uniform case of Theorem 16.

The non-uniform case of Theorem 16 follows the exact same logic; the only distinction is that we are no longer guaranteed that $T_n/\sqrt{n}$ has a limit, so we merely replace all instances of "$\lim_{n\to\infty} T_n/\sqrt{n}$" with either a "$\liminf_{n\to\infty}$" or a "$\limsup_{n\to\infty}$" depending on whether we are bounding from above or below. For example, the lower bound (7) becomes

$$\begin{aligned} T_n &\ge \frac{\mathsf{TSP}(P_1,\dots,P_n)}{\phi_0 + t\phi_0} \\ \implies \liminf_{n\to\infty} \frac{T_n}{\sqrt{n}} &\ge \lim_{n\to\infty} \frac{\mathsf{TSP}(P_1,\dots,P_n)}{\phi_0(1+t)\sqrt{n}} = \frac{\beta_{\text{TSP}} \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x}{\phi_0(1+t)} \ge \frac{0.6277 \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x}{\phi_0(1+t)} \\ \implies \phi_0\sqrt{\max\{1,t\}} \cdot \liminf_{n\to\infty} \frac{T_n}{\sqrt{n}} &\ge \frac{\phi_0 0.6277\sqrt{\max\{1,t\}}}{\phi_0(1+t)} \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x > 0.1368 \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x \text{ whenever } t < 19 \\ \implies \liminf_{n\to\infty} \frac{T_n}{\sqrt{n}} &\ge 0.1368c \text{ whenever } t < 19. \end{aligned}$$

The same reasoning is applied for the balanced-medians-derived lower bound for $t \ge 19$.

The upper bound that $\limsup T_n/\sqrt{n} \le 3.47c$ is also immediate; we already proved this for $t \ge 1$ in Lemma 14, and when $t < 1$, we again eschew the sidekicks altogether and use the truck:

$$\limsup_{n\to\infty} \frac{T_n}{\sqrt{n}} \le \frac{\beta_{\text{TSP}}}{\phi_0} \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x \le \frac{0.9204}{\phi_0} \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x < 3.47c\,,$$

which completes the proof. □

6 Remarks

Informally, Theorem 16 says that the time to completion of a sidekick routing problem satisfies

$$\text{Time With Sidekicks} \propto \frac{\sqrt{n}}{\sqrt{\phi_0 \max\{\phi_0, \phi_1 k\}}},$$

whereas when no sidekicks are present, the time without sidekicks is simply the duration of the TSP tour, which satisfies

$$\text{Time Without Sidekicks} = \frac{\text{TSP}}{\phi_0} \propto \frac{\sqrt{n}}{\phi_0}.$$

Thus, we claim that the amount of improvement due to using sidekicks is

$$\frac{\text{Time Without Sidekicks}}{\text{Time With Sidekicks}} \propto \frac{\sqrt{n}/\phi_0}{\sqrt{n}/\sqrt{\phi_0 \max\{\phi_0, \phi_1 k\}}} = \max\{1, \sqrt{\phi_1 k/\phi_0}\}. \tag{15}$$

We note that all of the above remarks hold in both the uniform and non-uniform cases because, as we have also seen in Theorem 16, the difference between these two cases merely amounts to multiplication by a factor of $\iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x$. In order to estimate the duration of a tour with sidekicks, we therefore propose the formula

$$\text{Time With Sidekicks} = c\left(\frac{\phi_1}{\phi_0}, k\right) \cdot \min\left\{1, \sqrt{\frac{\phi_0}{\phi_1 k}}\right\} \cdot (\text{Time Without Sidekicks}) \tag{16}$$

where $c(\phi_1/\phi_0, k)$ is a proportionality constant that depends on the ratio ϕ_1/ϕ_0 (as opposed to ϕ_0 and ϕ_1, which would be redundant by scaling) and k. The fact that $0.1368 < \beta_{\text{SK1}} < 3.47$ (and similarly for β_{SK2}) indicates that

$$0.192 \le c(\phi_1/\phi_0, k) \le 4.863$$

because

$$\text{Time Without Sidekicks} \sim \sqrt{n} \cdot \beta_{\text{TSP}} \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x$$

and

$$\text{Time With Sidekicks} \sim \sqrt{n} \cdot \frac{\beta_{\text{SK1}}}{\sqrt{\phi_0 \max\{\phi_0, \phi_1 k\}}} \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x$$

as $n \to \infty$, and therefore

$$\begin{aligned} c\left(\frac{\phi_1}{\phi_0}, k\right) &= \frac{\text{Time With Sidekicks}}{\min\left\{1, \sqrt{\frac{\phi_0}{\phi_1 k}}\right\} \cdot (\text{Time Without Sidekicks})} \\ &= \frac{\sqrt{n} \cdot \frac{\beta_{\text{SK1}}}{\sqrt{\phi_0 \max\{\phi_0, \phi_1 k\}}} \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x}{\min\left\{1, \sqrt{\frac{\phi_0}{\phi_1 k}}\right\} \cdot \sqrt{n} \cdot \frac{\beta_{\text{TSP}}}{\phi_0} \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x} = \frac{\beta_{\text{SK1}}}{\beta_{\text{TSP}}}, \end{aligned}$$

and the result follows by using the standard estimate $\beta_{\text{TSP}} \approx 0.7124$ [3]. We also know that $c(\phi_1/\phi_0, k) \to 1$ as $\phi_1/\phi_0 \to 0$ for k fixed, beccause sidekicks are no longer useful under this assumption. By applying

Remark 15, we find that $c(\phi_1/\phi_0, k) \leq 2.808$ when $\phi_1 \gg \phi_0$ by substitutiting the improved bound for β_{SK1}.

6.1 Battery life constraints

In real-world applications, one must contend with battery life constraints on the sidekicks, which is equivalent to restricting the total amount of distance that each sidekick traverses. Our upper and lower bounds from Sections 5.2 and 5.3 can be generalized to address this additional constraint. For brevity's sake, we will only consider the impact of battery constraints on problem (SK2).

6.1.1 A lower bound for (SK2) with a battery constraint

Suppose that battery life constraints are present, so that each sidekick can travel a distance of at most ω. We can bound problem (SK2) from below, subject to this new condition, by adding the constraint that

$$\sum_{i \in \mathcal{S}} \|x_i - p_i\| + \|p_i - y_i\| \leq k\omega\,. \tag{17}$$

This is a lower bound and not a feasible solution, because it imposes an aggregate constraint on the distance traversed by all sidekicks, as opposed to an individual constraint on each sidekick. This is because (SK2) does not single out individual vehicles, but is sufficient for our purposes. We have:

Lemma 17. *Let ℓ denote the total distance traversed by the truck in any feasible solution to (SK2), i.e.*

$$\ell = \sum_{j=2}^{2n+1} \|z_j - z_{j-1}\|\,,$$

subject to the additional constraint (17). We have

1. $\mathsf{TSP}(p_1, \ldots, p_n) \leq \ell + 2k\omega$
2. $\mathsf{BMed}(p_1, \ldots, p_n; d) \leq d\ell + k\omega/2$ *for all* $d \geq 2$.

Proof. The argument is identical to Lemma 12; the only difference is that is expressed in terms of the variables ℓ and ω. □

If we re-scale by setting $k\omega = a\sqrt{n}$ and $\ell = b\sqrt{n}$ and taking a limit as $n \to \infty$, then the above lemma says that the truck distance coefficient b (representing the scaled distance that the truck traverses) satisfies the lower bounds

1. $b \geq \beta_{\text{TSP}} - 2a$
2. $b \geq \sup_{d \geq 2, d \in \mathbb{Z}} \frac{\beta_{\mathsf{BMed}}(d) - a/2}{d}$

with probability one, as shown in Figure 8. Note that, per Remark 10, we see that (setting $c = 0.2886$)

$$b \geq \sup_{d \geq 100, d \in \mathbb{Z}} \frac{c\sqrt{d} - a/2}{d} \sim \left.\frac{c\sqrt{d} - a/2}{d}\right|_{d = a^2/c^2} = \frac{c^2}{2a}\,. \tag{18}$$

We will next derive an upper bound with similar relationships.

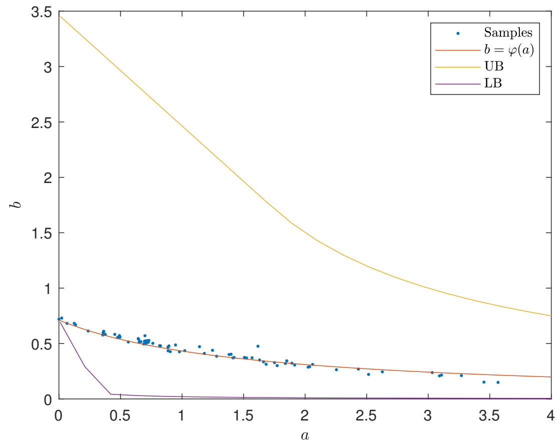

Figure 8: Upper and lower bounds for the truck length coefficient b, in terms of the sidekick bound coefficient a. The function $\varphi(a)$ lies within these bounds.

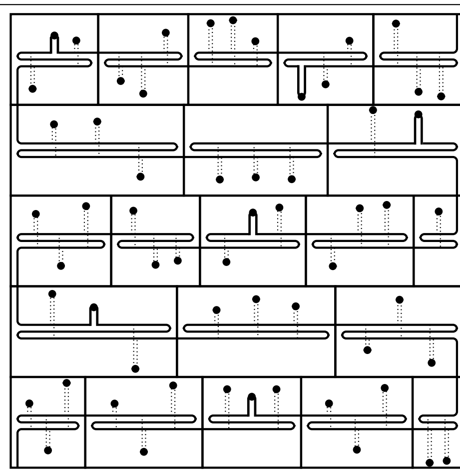

Figure 9: A zig-zag tour with certain customer points visited by the truck, such as the upper left cell and the center cell in the bottom (among others).

6.1.2 An upper bound for (SK2) with a battery constraint

Our upper bounding strategy closely follows the "zig-zag" strategy of Section 5.3. However, in order to accommodate the battery life constraint, we also allow for the possibility of the truck visiting certain customers directly via vertical trips, as shown in Figure 9. We construct a "zig-zag" tour as follows, ignoring rounding for notational convenience:

- If $k\omega \leq \sqrt{3n}$, then divide the region into rectangles of height $h = \sqrt{3/n}$. In addition to traversing these rectangles in the same way as in Section 5.3, also visit an arbitrary fraction $p = 1 - k\omega/\sqrt{3n}$ of customer points directly with the truck using vertical trips as in Figure 9. Up to an additive constant, the length ℓ of the truck's tour is at most

$$\frac{3}{h} + phn = 2\sqrt{3n} - k\omega$$

and by our selection of p, the aggregate length of all sidekick trips does not exceed $k\omega$.

- If $k\omega > \sqrt{3n}$, then divide the region into rectangles of height $h = k\omega/n$ and perform a standard zig-zag tour with no additional vertical trips made by the truck. Up to an additive constant, the length ℓ of the truck's tour is $n/(k\omega)$ and the aggregate length of all sidekick trips does not exceed $k\omega$.

We again re-scale by setting $k\omega = a\sqrt{n}$ and $\ell = b\sqrt{n}$. For large n, it is routine to construct feasible individual sidekick tours for the two strategies described above, using the fact that the aggregate lengths are at most $k\omega$ and that the vertical distances between customer points and the zig-zag tour follow a uniform distribution. Taking a limit as $n \to \infty$, the truck distance coefficient b satisfies the upper bounds

$$b \leq \begin{cases} \sqrt{12} - a & \text{if } a \leq \sqrt{3} \\ \frac{3}{a} & \text{otherwise} \end{cases}$$

as shown in Figure 8.

6.1.3 Asymptotic behavior

Regarded as a function $b = \varphi(a)$, the upper and lower bounds that we just established are both of the form

$$\varphi(a) = \begin{cases} \xi - \eta a & \text{for } a \leq \bar{a} \\ \frac{\zeta}{a} & \text{otherwise} \end{cases}$$

for positive coefficients ξ, η, and ζ and a threshold $\bar{a}$, and are within a constant factor of one another. Based on the preceding analysis and some geometric intuition, we propose the following approximation:

$$\varphi(a) = \frac{\zeta_0}{a + \zeta_0/\beta_{\text{TSP}}}\,, \tag{19}$$

with $\zeta_0 = 1.10$ and $\beta_{\text{TSP}} \approx 0.7124$; see [3] for the latter. Figure 8 shows that $\varphi(a)$ lies within our bounds. We obtained the value of ζ_0 by numerical simulations using Google OR-tools, by solving the following simplification of (SK2) that disregards the temporal interaction between sidekicks and the truck, and merely minimizes the length of the truck tour subject to a constraint on the distance traversed by the sidekicks:

$$\begin{aligned} \underset{\boldsymbol{x}}{\text{minimize}} \quad & \mathsf{TSP}(x_1, \ldots, x_n) \quad s.t. \\ & \sum_{i=1}^{n} \|x_i - p_i\| \leq k\omega\,. \end{aligned} \tag{SK-simple}$$

This estimation process is as follows: fix $n = 100$, and solve (SK-simple) for various values of $k\omega$ and points p_i sampled uniformly at random in the unit square. For each value of $k\omega$, let ℓ^* denote the objective value, so that we can estimate $\hat{a} = k\omega/\sqrt{n}$ and $\hat{b} = \ell^*/\sqrt{n}$. Given a large collection of such estimates, we estimate ζ_0 via standard regression techniques.

6.1.4 Approximation formulas with battery life constraints

The analysis in sections 6.1.1-6.1.3 characterizes a Pareto frontier between the distance ℓ traversed by the truck and the cumulative distance traversed by all the sidekicks. Given speeds ϕ_0, ϕ_1 and a constraint that each of the k sidekicks can travel a total distance ω, we see that the total makespan is at least $\max\{\ell/\phi_0, k\omega/\phi_1\}$, so a valid lower bounding continuous approximation formula relating battery life and makespan is the solution to the problem

$$\begin{aligned} \underset{s,t}{\text{minimize}} \max\left\{\frac{s}{k\phi_1}, \frac{t}{\phi_0}\right\} \quad & s.t. \\ s &\leq a \\ t &\geq \varphi(s)\,, \end{aligned} \tag{20}$$

where s represents sidekick travel distance and t represents truck travel distance, and both have been scaled by $\sqrt{n}$ as we have done previously. It is easy to see that the solution (s^*, t^*) to this problem always satisfies $t^* = \varphi(s^*)$, and so, when the battery life constraint is binding, the optimal solution has $s^* = a$ and $t^* = \varphi(a)$, resulting in a makespan that is at least

$$\frac{t^*}{\phi_0} = \frac{\varphi(a)}{\phi_0} = \frac{\zeta_0}{\phi_0(a + \zeta_0/\beta_{\text{TSP}})}\,.$$

Thus, when each sidekick is constrained to traverse a distance of at most ω, then when the battery constraint is binding, we find (by substituting $a = k\omega/\sqrt{n}$ and $\ell^* = t^*\sqrt{n}$) that the makespan T_n must be at least

$$T_n \gtrsim \frac{\zeta_0 n}{\phi_0(k\omega + \zeta_0\sqrt{n}/\beta_{\text{TSP}})} \tag{21}$$

for large n.

Note that the objective function of (21) was derived by bounding the makespan from below, because it is merely the maximum of the truck's time and the average time of the sidekicks.. It is straightforward to derive an objective function based on bounding the makespan above, with the same proportionality, by using the zig-zag argument from Section 6.1.2; we omit it here for the sake of brevity.

6.2 A budget constraint

This paper has thus far been concerned with minimizing time to completion, although it is also a challenging and important problem to minimize overall costs; see for instance [28]. We can accomplish this using much of the same machinery as in Section 6.1; the only distinction is that Section 6.1 imposes a hard constraint on the distance covered by sidekicks, whereas cost minimization penalizes both distances at different rates. In particular, we use the same relationship φ established in (19), although now the counterpart to problem (20) becomes

$$\begin{aligned} \underset{s,t}{\text{minimize}} \max\left\{\frac{s}{k\phi_1}, \frac{t}{\phi_0}\right\} \quad s.t. \\ \alpha_0 t + \alpha_1 s \le \xi \\ t \ge \varphi(s) \; ; \end{aligned} \tag{22}$$

the only change is that we have replaced the drone battery life constraint $s \le a$ with a linear constraint, where α_0 and α_1 are unit costs for the truck and sidekicks respectively, and ξ is an overall budget. Again, we find that $t^* = \varphi(s^*)$ at optimality, and so when the budget constraint is binding, we have $\alpha_0\varphi(s^*) + \alpha_1 s^* = \xi$. This is solvable as a quadratic in s^*, and the resultant (scaled) makespan t/ϕ_0 turns out to be

$$\frac{1}{\phi_0} \cdot \frac{2\alpha_1\beta_{\text{TSP}}\zeta_0}{\beta_{\text{TSP}}\xi + \alpha_1\zeta_0 - \sqrt{\beta_{\text{TSP}}^2\xi^2 + \alpha_1^2\zeta_0^2 - 2(2\,\alpha_0\alpha_1\beta_{\text{TSP}}^2 - \alpha_1\beta_{\text{TSP}}\xi)\zeta_0}} \, .$$

6.3 Launch and retrieval times

A further important complication is to model the amount of time needed to launch or retrieve a sidekick. If each sidekick launch and retrieval requires a fixed duration δ, then it is routine to verify that both of the lower bounds from Section 5.2 increase by at least $2\delta n$; this is because the relevant inequalities from Lemma 12 merely become

1. $\mathsf{TSP}(p_1, \ldots, p_n) \le (\phi_0 + k\phi_1)T_n - 2\delta n$
2. $\mathsf{BMed}(p_1, \ldots, p_n; d) \le (d\phi_0 + k\phi_1)T_n - 2\delta n$ for all $d \ge 2$.

by carrying the proof out under this assumption. The zig-zag upper bound of Section 5.3 also increases by at most $2\delta n$; this is because the amount of time to service a rectangle with width w and height h is now at most $3w/\phi_0 + h/\phi_1 + 2\delta k$, as opposed to only $3w/\phi_0 + h/\phi_1$ as in the proof of Lemma 13.

Note, however, that although the lower and upper bounds are affected additively in the same way (by adding $2\delta n$), the proportionality is now different, because Theorem 16 says that the makespan is proportional to $\sqrt{n}$, and fixed costs are (for obvious reasons) linear in n. Thus, in order to extend Theorem 16 to this case, it becomes necessary to assume that $\delta \sim 1/\sqrt{n}$. In particular, we now find that if $\delta = \tau/\sqrt{n}$ for some constant τ, then Theorem 16 says that

$$\frac{T_n}{\sqrt{n}} \to \frac{\beta_{\text{SK1}}}{\sqrt{\phi_0 \max\{\phi_0, \phi_1 k\}}} + 2\tau \, .$$

7 Computational results

This section describes two computational experiments. The first consists of uniform samples in the unit square with Euclidean distances, and the second consists of samples taken around the Los Angeles Metropolitan area with respect to a road network. For both experiments, we sampled $n = 100$ points

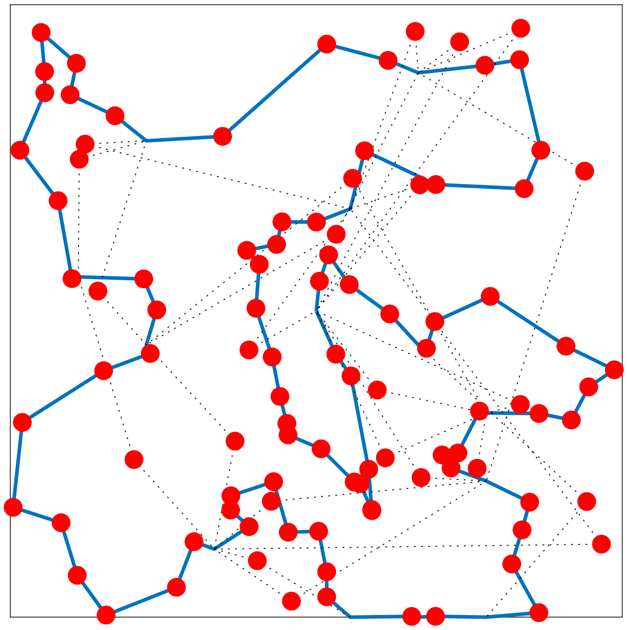

Figure 10: A solution to problem (SK2) as obtained via the LKH-3 heuristic, for $\phi_1 = 0.5$ and $k = 5$.

and considered $k \in \{1, 2, 3, 4, 5, 6\}$, and conducted 10 experiments for each scenario. Both sections consider problem (SK2), so that the truck is permitted to visit customer demand points.

In order to solve the problem instances, we implemented a penalty function for the LKH-3 heuristic solver, as has been done previously for many variants of VRP [12]; see Section C of the Online Supplement. As our problem allows for continuous placement of launch sites (the variables $\boldsymbol{x}$, $\boldsymbol{y}$, and $\boldsymbol{z}$ in the original formulation (SK2)), but LKH-3 requires discrete inputs, for both problem instances we discretized the set of possible launch sites into a 12×12 grid. The goal of this section is to estimate the proportionality constant $c(\phi_1/\phi_0, k)$, as expressed in equation (16); recall that we already established that $0.192 \leq c(\phi_1/\phi_0, k) \leq 4.863$ in Section 6, so the experiments will ideally tighten these bounds (non-rigorously) and give some insight into the extent to which $c(\phi_1/\phi_0, k)$ varies.

7.1 Uniformly distributed demand with Euclidean travel

In our first experiment, we sampled $n = 100$ points in the unit square plus a "depot" centered at $(0.5, 0.5)$, and assume $\phi_0 = 1$ without loss of generality. We computed the TSP tour of these demand points with Euclidean distances, which we then compare with the solution to problem (SK2) as determined from LKH-3; see Figure 10 for an example. The objective values for the sidekick tours, and the resulting estimates of $c(\phi_1/\phi_0, k)$, are shown in Tables 1 and 2.

Overall, the ratios shown in Table 2 indicate that $0.9 \leq c(\phi_1/\phi_0, k) \leq 2$, with $c(\phi_1/\phi_0, k)$ increasing in k and ϕ_1. The entries where $c(\phi_1/\phi_0, k) < 1$ are unsurprising, as they mostly occur when $\phi_0/(\phi_1 k) < 1$, in which case the entry $\min\{1, \sqrt{\phi_0/(\phi_1 k)}\}$ in (16) is simply 1, and the prediction is that the tour duration is simply

$$c\left(\frac{\phi_1}{\phi_0}, k\right) \cdot (\text{Time Without Sidekicks}) \,.$$

Objective values of (SK2)

ϕ_1 =0.3	7.364	7.184	7.049	6.736	6.728	6.561
0.5	7.226	6.938	6.923	6.641	6.326	6.132
0.875	6.880	6.544	6.083	5.72	5.702	5.642
1.25	6.541	5.996	5.495	4.890	4.776	4.476
1.625	6.199	5.650	5.17	4.604	4.224	4.222
2	6.152	5.629	4.927	4.333	4.082	4.06
3	6.121	5.442	4.838	4.201	3.895	3.54
	$k = 1$	2	3	4	5	6

Table 1: Makespans of sidekick-assisted TSPs (i.e. the optimal objective value to problem (SK2)), for $\phi_0 = 1$ and varying values of ϕ_1 and k.

Estimates of $c(\phi_1/\phi_0, k)$

ϕ_1 =0.3	0.9728	0.9497	0.9311	0.9756	1.0889	1.1632
0.5	0.9551	0.9166	1.1206	1.2412	1.3219	1.4035
0.875	0.9091	1.1442	1.3024	1.4142	1.5768	1.7085
1.25	0.9662	1.2528	1.4061	1.4449	1.5780	1.6202
1.625	1.0443	1.3461	1.5088	1.5513	1.5911	1.7422
2	1.1498	1.4877	1.5949	1.6194	1.7059	1.8587
3	1.4011	1.7615	1.9179	1.9231	1.9936	1.9845
	$k = 1$	2	3	4	5	6

Table 2: Estimates of $c(\phi_1/\phi_0, k)$, determined by comparing the values in Table 1 to the length of a standard TSP tour and the correction factor $\min\{1, \sqrt{\phi_0/(\phi_1 k)}\}$ from (16).

Objective values of (SK2) (in minutes)

	457	440	424	406	404	380
	$k = 1$	2	3	4	5	6

Table 3: Makespans of sidekick-assisted TSPs (i.e. the optimal objective value to problem (SK2)), for $\phi_0 = 1$ and varying values of ϕ_1 and k.

When $\phi_0/(\phi_1 k)$ is very small, the sidekicks provide minimal benefit and the problem effectively reduces to a standard TSP. However, in the transitional region where $\phi_0/(\phi_1 k)$ is moderately below 1, sidekicks can still provide some benefit, though this improvement is not fully captured by our current analytical framework. This explains why we observe ratios below 1 in these cases, but we would expect these ratios to approach 1 as the sidekicks become progressively slower.

7.2 A road network

In our second experiment, we sampled $n = 100$ points in the Los Angeles metropolitan region plus a "depot" centered at the University of Southern California. Since the focus of this paper has been on the case where the sidekicks are slower than the truck, we assume that the sidekick speed is equal to the *bicycling* speed from one point to the next, as opposed to driving; see Figure 11 for an example of a solution. The process for modelling the difference in speeds requires more nuance, as none of the vehicles travels at a constant speed. We therefore estimate the duration of the sidekick tour as follows:

1. Compute the duration T_1 of a TSP tour of the n points with respect to driving time on a road network, as computed via a locally-hosted OSRM server [18].
2. Compute the duration T_2 of the same tour as in step 1 with respect to bicycling time, again using OSRM. The ratio between speeds, ϕ_1/ϕ_0, is T_1/T_2.
3. Compute $c(\phi_1/\phi_0, k)$ according to (16).

Figure 11 shows an example of a solution determined with LKH-3. Our experiments found that $\phi_1/\phi_0 = 0.346$, and the results are shown in Table 3. The best comparison between these results and those from the preceding section are obtained by comparing them to the top row of Table 1, in which $\phi_1/\phi_0 = 0.3$ and $0.9728 \leq \phi_1/\phi_0 \leq 1.1632$, as opposed to $\phi_1/\phi_0 = 0.346$ and $0.9690 \leq \phi_1/\phi_0 \leq 1.1896$ in our present analysis. The consistency of these values across both experimental settings offers some validation for our theoretical framework and suggests that the proportionality constant exhibits similar behavior regardless of whether travel occurs in Euclidean space or on a road network.

8 Conclusions

We have studied the limiting behavior of sidekick-assisted routing problems in the Euclidean plane and found that the improvements introduced by adding sidekicks can be predicted based on the relationship between $\sqrt{\phi_0}$ and $\sqrt{\phi_1 k}$. Our model at present does not fully capture the improvements

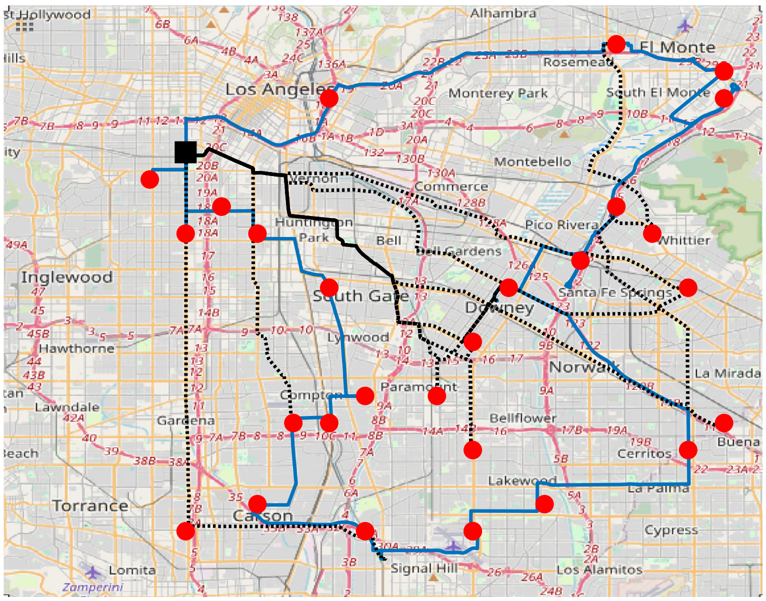

Figure 11: A solution to problem (SK2) as obtained via the LKH-3 heuristic for demand locations in Los Angeles with $k = 3$ sidekicks, whose speed is equivalent to that of a bicycle. Dashed lines correspond to sidekick routes and the blue line indicates the vehicle tour. Although our experiments used $n = 100$, this picture uses $n = 30$ for visual clarity.

Estimates of $c(\phi_1/\phi_0, k)$

	0.9690	0.9320	0.9398	1.0390	1.1566	1.1896
	$k = 1$	2	3	4	5	6

Table 4: Estimates of $c(\phi_1/\phi_0, k)$, determined by comparing the values in Table 3 to the length of a standard TSP tour and the correction factor $\min\{1, \sqrt{\phi_0/(\phi_1 k)}\}$ from (16).

that are, somewhat surprisingly, realized when k and ϕ_1 are simultaneously both small, but is within a sufficiently small margin of error to still serve a useful purpose. The main contributions of this paper are threefold: First, we provide a rigorous asymptotic analysis of sidekick routing problems with multiple sidekicks, establishing upper and lower bounds that characterize the fundamental scaling behavior of the problem. Second, we demonstrate that sidekick systems can be beneficial even when the sidekicks are slower than the host vehicle, provided there are sufficiently many of them, which contradicts the conventional wisdom that sidekicks must be faster to be useful. Third, we develop a practical formula (16), with an empirically validated proportionality constant $c(\phi_1/\phi_0, k)$ estimated in Section 7, that allows practitioners to estimate the potential benefits of adopting a sidekick system.

There remain many open questions: for example, what happens when sidekicks are able to visit more than one customer node before returning to the truck? What happens when the truck is itself capacitated and must make returns to the depot? What happens when sidekick battery life considerations come into play? Additional promising research directions include tightening the theoretical bounds through space-filling curve approaches [4, 6], investigating heterogeneous fleets with varying sidekick speeds, and analyzing the impact of time-dependent travel speeds on routing decisions. Because of the slow rate of adoption of such systems in real-world deployment – likely due to practical considerations, as discussed in the introduction of this paper – it is difficult at present to determine the full extent to which our predictions hold, but we believe this work provides a solid foundation for future algorithmic and theoretical developments.

References

[1] 118th Congress. Faa reauthorization act of 2024. Public Law No. 118-63, 5 2024. H.R.3935.

[2] Niels Agatz, Paul Bouman, and Marie Schmidt. Optimization approaches for the traveling salesman problem with drone. *Transportation Science*, 52(4):965–981, 2018.

[3] D. Applegate, W. Cook, D. S. Johnson, and N. J. A. Sloane. Using large-scale computation to estimate the Beardwood-Halton-Hammersley TSP constant. Presentation at 42 SBPO, 2010.

[4] John J Bartholdi III and Loren K Platzman. Heuristics based on spacefilling curves for combinatorial problems in euclidean space. *Management Science*, 34(3):291–305, 1988.

[5] J. Beardwood, J. H. Halton, and J. M. Hammersley. The shortest path through many points. *Mathematical Proceedings of the Cambridge Philosophical Society*, 55(4):299–327, 1959.

[6] Dimitris Bertsimas and Michelangelo Grigni. Worst-case examples for the spacefilling curve heuristic for the euclidean traveling salesman problem. *Operations Research Letters*, 8(5):241–244, 1989.

[7] James Campbell, Donald Sweeney, and Juan Zhang. Strategic design for delivery with trucks and drones, 04 2017. Supply Chain & Analytics Report SCMA-2017-0201.

[8] John Gunnar Carlsson and Siyuan Song. Coordinated logistics with a truck and a drone. *Management Science*, 64(9):4052–4069, 2017.

[9] Sung Hoon Chung, Bhawesh Sah, and Jinkun Lee. Optimization for drone and drone-truck combined operations: A review of the state of the art and future directions. *Computers & Operations Research*, 123:105004, 2020.

[10] Zac Dukowitz. Amazon files new patent to use drones and trucks together for deliveries. `https://uavcoach.com/amazon-drones-trucks/`, 2021. UAVCoach.com.

[11] Julia Gaudio and Patrick Jaillet. An improved lower bound for the traveling salesman constant. *Operations Research Letters*, 48(1):67–70, 2020.

[12] Keld Helsgaun. An extension of the lin-kernighan-helsgaun tsp solver for constrained traveling salesman and vehicle routing problems. *Roskilde: Roskilde University*, 12:966–980, 2017.

[13] Ines Khoufi, Anis Laouiti, and Cedric Adjih. A survey of recent extended variants of the traveling salesman and vehicle routing problems for unmanned aerial vehicles. *Drones*, 3(3):66, 2019.

[14] J.K. Korpela. *Mathematical Expressions*. Suomen E-painos Oy, 2014.

[15] Giusy Macrina, Luigi Di Puglia Pugliese, Francesca Guerriero, and Gilbert Laporte. Drone-aided routing: A literature review. *Transportation Research Part C: Emerging Technologies*, 120:102762, 2020.

[16] K McGivney and JE Yukich. Asymptotics for geometric location problems over random samples. *Advances in Applied Probability*, 31(3):632–642, 1999.

[17] Mohammad Moshref-Javadi and Matthias Winkenbach. Applications and research avenues for drone-based models in logistics: A classification and review. *Expert Systems with Applications*, 177:114854, 2021.

[18] OpenStreetMap contributors. Planet dump retrieved from https://planet.osm.org . `https://www.openstreetmap.org`, 2023.

[19] Alena Otto, Niels Agatz, James Campbell, Bruce Golden, and Erwin Pesch. Optimization approaches for civil applications of unmanned aerial vehicles (uavs) or aerial drones: A survey. *Networks*, 72(4):411–458, 2018.

[20] Stefan Poikonen and James F Campbell. Future directions in drone routing research. *Networks*, 77(1):116–126, 2021.

[21] Stefan Poikonen, Xingyin Wang, and Bruce Golden. The vehicle routing problem with drones: Extended models and connections. *Networks*, 70(1):34–43, 2017.

[22] 6 River Systems. Meet Chuck. `https://6river.com/meet-chuck/`, 2023. 6river.com.

[23] Daniela Rojas Viloria, Elyn L Solano-Charris, Andrés Muñoz-Villamizar, and Jairo R Montoya-Torres. Unmanned aerial vehicles/drones in vehicle routing problems: a literature review. *International Transactions in Operational Research*, 28(4):1626–1657, 2021.

[24] Yogesh Shinde. Drone package delivery market to reach at a usd 33,813.2 million by 2033. `https://scoop.market.us/drone-package-delivery-market-news/`, 2024. Scoop.market.us.

[25] J. M. Steele. Subadditive Euclidean functionals and nonlinear growth in geometric probability. *The Annals of Probability*, 9(3):pp. 365–376, 1981.

[26] J.M. Steele. *Probability Theory and Combinatorial Optimization*. CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics, 1987.

[27] Jack Stewart. A drone-slinging UPS van delivers the future. `https://www.wired.com/2017/02/drone-slinging-ups-van-delivers-future/`, 2017. WIRED.

[28] Gizem Ozbaygin Tinic, Oya E Karasan, Bahar Y Kara, James F Campbell, and Aysu Ozel. Exact solution approaches for the minimum total cost traveling salesman problem with multiple drones. *Transportation Research Part B: Methodological*, 168:81–123, 2023.

[29] Transport Canada. Regulations amending the canadian aviation regulations (rpas-beyond visual line-of-sight and other operations). Canada Gazette, Part II, 3 2025. Effective April 1, 2025, with full implementation by November 4, 2025.

[30] Xingyin Wang, Stefan Poikonen, and Bruce Golden. The vehicle routing problem with drones: Several worst-case results. *Optimization Letters*, 11, 04 2016.

[31] J.E. Yukich. *Probability Theory of Classical Euclidean Optimization Problems*. Lecture Notes in Mathematics. Springer Berlin Heidelberg, 2006.

[32] Juan Zhang. *Economic and Environmental Impacts of Drone Delivery*. PhD thesis, University of Missouri-Saint Louis, 2021.

Online supplement to "Coordinated logistics with a truck and multiple sidekicks"

A Proof of Theorem 9

The following three lemmas are textbook-level results that we state without proof:

Lemma 18 (Stirling's approximation)**.** *The gamma function* $\Gamma(x)$ *satisfies*

$$\log \Gamma(x+1) = x \log x - x + \frac{1}{2}\log x + \frac{1}{2}\log 2 + \frac{1}{2}\log \pi + \mathcal{O}(1/x).$$

Lemma 19. *Let* $f : \mathbb{R} \mapsto \mathbb{R}$ *be a real-valued function and let* $\mathcal{B}_d(r) \subset \mathbb{R}^d$ *be the ball of radius* r *centered about the origin. We have*

$$\int_{\mathcal{B}_d(r)} f(\|x\|)dx = \int_0^r S_{d-1}(t)f(t)dt,$$

where $S_{d-1}(t)$ *is the surface area of a* $(d-1)$*-sphere of radius* t*, which is given by*

$$S_{d-1}(t) = \frac{2\pi^{d/2}}{\Gamma(d/2)} t^{d-1}.$$

Lemma 20. *The volume of a* d*-dimensional ball of radius* r *is* $\pi^{d/2}r^d/\Gamma(d/2+1)$.

Lemma 21. *Let* $l > 0$ *and let* $\mathcal{D} \subset \mathbb{R}^{2n}$ *denote the set of all* n*-tuples* $(u_1, \ldots, u_n)$ *of points in* $\mathbb{R}^2$ *such that*

$$\sum_{i=1}^n \|u_i\| < l.$$

The volume of $\mathcal{D}$*,* $\mathrm{Vol}(\mathcal{D})$*, satisfies*

$$\mathrm{Vol}(\mathcal{D}) = \frac{(2\pi)^n}{\Gamma(2n+1)} \cdot l^{2n}.$$

Proof. This is just the integral

$$\int_{\mathcal{B}_2(\ell)} \int_{\mathcal{B}_2(\ell - \|\mathbf{u}_n\|)} \cdots \int_{\mathcal{B}_2(\ell - \sum_{i=3}^n \|\mathbf{u}_i\|)} \int_{\mathcal{B}_2(\ell - \sum_{i=2}^n \|\mathbf{u}_i\|)} 1 \, d\mathbf{u}_1 \, d\mathbf{u}_2 \, \cdots \, d\mathbf{u}_{n-1} \, d\mathbf{u}_n \, ,$$

which we can compute by induction. For $n = 1$,

$$\int_{\mathcal{B}_2(\ell)} 1 du_1 = \pi \ell^2 = \frac{(2\pi)^1}{\Gamma(2(1)+1)} \ell^{2(1)}. \qquad \text{(Lemma 20)}$$

Suppose the relation holds for all ℓ' for the set of all $(n-1)$-tuples such that $\sum_{i=1}^{n-1} \|\mathbf{u}_i\| \leq \ell'$. Then if $\mathcal{D}$ is the set of all n-tuples $(\mathbf{u}_1, \ldots, \mathbf{u}_n)$ of points in $\mathbb{R}^2$ such that $\sum_{i=1}^n \|\mathbf{u}_i\| \leq \ell$, we have

$$\begin{aligned}
\mathrm{Vol}(\mathcal{D}) &= \int_{\mathcal{B}_2(\ell)} \int_{\mathcal{B}_2(\ell - \|\mathbf{u}_n\|)} \cdots \int_{\mathcal{B}_2(\ell - \sum_{i=3}^n \|\mathbf{u}_i\|)} \int_{\mathcal{B}_2(\ell - \sum_{i=2}^n \|\mathbf{u}_i\|)} 1 \, d\mathbf{u}_1 \, d\mathbf{u}_2 \cdots d\mathbf{u}_{n-1} \, d\mathbf{u}_n \\
&= \int_{\mathcal{B}_2(\ell)} \frac{(2\pi)^{n-1}}{\Gamma(2(n-1)+1)} (\ell - ||u_n||)^{2(n-1)} du_n && \text{(induction hypothesis)} \\
&= \int_0^\ell 2\pi t \frac{(2\pi)^{n-1}}{\Gamma(2(n-1)+1)} (\ell - t)^{2(n-1)} dt && \text{(Lemma 19)} \\
&= \frac{(2\pi)^n}{\Gamma(2(n-1)+1)} \int_0^\ell t(\ell - t)^{2(n-1)} dt \\
&= \frac{(2\pi)^n}{\Gamma(2n-1)} \cdot \frac{\ell^{2n}}{2n(2n-1)} \\
&= \frac{(2\pi)^n}{\Gamma(2n+1)} \ell^{2n}.
\end{aligned}$$

□

We are now ready to prove Theorem 9:

Proof of theorem 9. As $\beta_{\mathsf{BMed}}(d)$ is independent of the demand distribution, we can arrive at a lower bound by first assuming we are in the case that the X_i are i.i.d. $\mathsf{Unif}([0,1]^2)$. We employ the union bound.

$$\begin{aligned}
\mathbb{P}(\mathrm{BMed}(X_1, \ldots, X_n; d) < l) &= \mathbb{P}(\text{some selection of, and assignment to, medians is of cost} < l) \\
&\leq \text{sum over all selections and assignments of } \mathbb{P}(\text{cost of selection, assignment} < l) \\
&= (\text{\# ways select medians})(\text{\# ways assign points})\mathbb{P}(\text{cost of arbitrary choice} < l) \\
&= \binom{n}{p} \frac{(n-p)!}{(d!)^p} \mathbb{P}(\text{cost of arbitrary selection and assignment} < l),
\end{aligned}$$

where for our asymptotic results we are justified in disregarding the ceiling and assuming $d+1$ divides n. To obtain an upper bound on the probability that an arbitrary selection and assignment has cost less than l we fix our median indices, $\mathcal{S}$, and our assignment map μ and first recall that $X_{\mu(i)}$ is the median assigned to point X_i. Because we can reorder and adjust μ and $\mathcal{S}$ accordingly, we may assume without loss of generality that the X_i are ordered such that the medians we have selected are the last p points, $X_{n-p+1}, \ldots, X_n$. We then have the cost of a particular selection and assignment is given by $\sum_{i=1}^{n-p} \|X_i - X_{\mu(i)}\|$. Let

$$\mathcal{E}(l) := \left\{ x_1, \ldots, x_{n-p} \in \mathbb{R}^2 : \sum_{i=1}^{n-p} \|x_i - x_{\mu(i)}\| < l \right\}.$$

Then recalling that the X_i are drawn uniformly from the unit square we have,

$$\begin{aligned}\mathbb{P}\left(\sum_{i=1}^{n-p}\|X_i - X_{\mu(i)}\| < l\right) &= \mathbb{P}((X_1,\dots,X_{n-p}) \in \mathcal{E}(l)) \\ &= \mathrm{Vol}(\mathcal{E}(l) \cap [0,1]^2) \\ &\le \mathrm{Vol}(\mathcal{E}(l)).\end{aligned}$$

To compute this volume we make the volume preserving transformation $u_i := x_i - x_{\mu(i)}$ (that is, translate each median point to the origin and move its assigned set commensurately) and consider

$$\mathcal{E}'(l) := \left\{u_1,\dots,u_{n-p} \in \mathbb{R}^2 : \sum_{i=1}^{n-p}\|u_i\| < l\right\}.$$

Then $\mathrm{Vol}(\mathcal{E}(l)) = \mathrm{Vol}(\mathcal{E}'(l))$. By Lemma 21 we have

$$\mathrm{Vol}(\mathcal{E}'(l)) = \frac{(2\pi)^{n-p}}{\Gamma(2(n-p)+1)} \cdot l^{2(n-p)}.$$

Thus for all selections of $\mathcal{S}$ and μ,

$$\begin{aligned}\mathbb{P}\left(\sum_{i=1}^{n-p}\|X_i - X_{\mu(i)}\| < l\right) &\le \mathrm{Vol}(\mathcal{E}'(l)) \\ &= \frac{(2\pi)^{n-p}}{\Gamma(2(n-p)+1)} \cdot l^{2(n-p)}.\end{aligned}$$

Combining the above

$$\mathbb{P}(\mathrm{BMed}(X_1,\dots,X_n;d) < l) \le \binom{n}{p}\frac{(n-p)!}{(d!)^p} \cdot \frac{(2\pi)^{n-p}}{\Gamma(2(n-p)+1)} \cdot l^{2(n-p)}.$$

Taking logarithms then yields

$$\begin{aligned}\log \mathbb{P}(\mathrm{BMed}(X_1,\dots,X_n;d) < l) &\le \log\left(\binom{n}{p}\frac{(n-p)!}{(d!)^p} \cdot \frac{(2\pi)^{n-p}}{\Gamma(2(n-p)+1)} \cdot (l)^{2(n-p)}\right) \\ &= \log\left(\frac{\Gamma(n+1)}{\Gamma(p+1)\Gamma(n-p+1)}\frac{\Gamma(n-p+1)}{\Gamma(d+1)^p} \cdot \frac{(2\pi)^{n-p}}{\Gamma(2(n-p)+1)} \cdot l^{2(n-p)}\right) \\ &= \Bigg[\left(1-(d+1)^{-1}\right)\log(2) + \left(1-(d+1)^{-1}\right)\log(\pi) \\ &\quad + \left(2-2(d+1)^{-1}\right)\log(l) + \log(n) + 1 - \frac{1}{d+1}\log\left(\frac{n}{d+1}\right) \\ &\quad - (d+1)^{-1} - \frac{\log(d!)}{d+1} - \left(2-2(d+1)^{-1}\right)\log\left(2n - 2\frac{n}{d+1}\right)\Bigg]n \\ &\quad - 1/2\log(2) - 1/2\log(\pi) + \mathcal{O}(1/n),\end{aligned}$$

where we have employed Lemma 18. It is clear that this upper bound goes to negative infinity as n goes to infinity if and only if the coefficient of n is negative. We then have as $n \to \infty$,

$$\begin{aligned}
&\mathbb{P}(\mathrm{BMed}(X_1, \ldots, X_n; d) < l) \to 0 \\
&\Uparrow \\
&\left(1 - (d+1)^{-1}\right) \log(2) + \left(1 - (d+1)^{-1}\right) \log(\pi) \\
&+ \left(2 - 2(d+1)^{-1}\right) \log(l) + \log(n) + 1 - \frac{1}{d+1} \log\left(\frac{n}{d+1}\right) \\
&- (d+1)^{-1} - \frac{\log(d!)}{d+1} - \left(2 - 2(d+1)^{-1}\right) \log\left(2n - 2\frac{n}{d+1}\right) < 0 \\
&\Updownarrow \\
&l < \frac{\sqrt{2}d(d!)^{1/(2d)}}{\sqrt{\pi e}(d+1)^{(2d+1)/(2d)}} \cdot \sqrt{n}.
\end{aligned}$$

The result then follows easily from the almost sure convergence to β_{BMed}. □

B Proof of Lemma 14

If f is absolutely continuous, then it can be approximated arbitrarily well with finitely many step functions on $\mathcal{R}$, and we therefore assume without loss of generality that f takes precisely this form. To be more specific, we assume that $f(x) = \sum_{i=1}^m f_i \delta_i(x)$, where $\delta_i(x)$ is an indicator function representing membership in a square grid cell i. Let ϵ denote the area of each grid cell, so that $\iint_{\mathcal{R}} f(x)\,\mathrm{d}x = \sum_{i=1}^m \epsilon f_i = 1$, and let N_i denote the number of samples of $\{P_1, \ldots, P_n\}$ that belong to cell i (so that $\sum_{i=1}^m N_i = n$).

It is clear that we can construct a feasible tour by applying Lemma 13 to each grid cell and then "stitching" the tours within each grid cell together. Certainly, the amount of time needed to visit all N_i points in grid cell i is at most

$$\sqrt{\epsilon}\left(\frac{2\sqrt{3}}{\sqrt{\phi_0\phi_1 k}} \cdot \sqrt{N_i} + C_i\right)$$

for some constant C_i, and the amount of additional time needed to "stitch" all of the tours together is a constant C_0 that does not depend on n. Summing all of these together and letting $C = \sum_{i=0}^m C_i$, we have

$$\begin{aligned}
T(P_1, \ldots, P_n) &\le \sqrt{\epsilon}\sum_{i=1}^m \left(\frac{2\sqrt{3}}{\sqrt{\phi_0\phi_1 k}} \cdot \sqrt{N_i} + C_i\right) + C_0 \\
\implies \frac{T(P_1, \ldots, P_n)}{\sqrt{n}} &\le \sqrt{\epsilon}\sum_{i=1}^m \left(\frac{2\sqrt{3}}{\sqrt{\phi_0\phi_1 k}} \cdot \sqrt{\frac{N_i}{n}}\right) + \frac{C}{\sqrt{n}}
\end{aligned}$$

and since $N_i/n \to \epsilon f_i$ with probability one, we see that

$$\limsup_{n\to\infty} \frac{T(P_1,\ldots,P_n)}{\sqrt{n}} \leq \sqrt{\epsilon}\sum_{i=1}^{m} \frac{2\sqrt{3}}{\sqrt{\phi_0\phi_1 k}} \cdot \sqrt{\epsilon f_i} = \frac{2\sqrt{3}}{\sqrt{\phi_0\phi_1 k}} \sum_{i=1}^{m} \epsilon\sqrt{f_i} = \frac{2\sqrt{3}}{\sqrt{\phi_0\phi_1 k}} \iint_{\mathcal{R}} \sqrt{f(x)}\,\mathrm{d}x$$

as desired.

C Penalty function for LKH-3 for sidekick routing

The code below replaces the file `Penalty_1_PDTSP.c` in LKH-3; we use the built-in attribute `Node->Color` to identify different customer nodes; a customer is "satisfied" if a node with their corresponding color has been visited twice (a launch and a retrieval), or if the customer node has been visited directly by the vehicle.

```
#include "LKH.h"
#include "Segment.h"
#define COLOR_COUNT 100

#define Pen1 100000 // 1000000 no 'revisiting a launch or retrieval node'
#define Pen2 100000 // 10000 drone number for each route
#define Pen3 100000 // 10000 drone number for the system

GainType Penalty_1_PDTSP()
{
    Node *N, *NextN, *instantneN, *trueneN, *prevN, *prevN2, *prevN3;
    GainType P = 0;
    GainType CurrentTime = 0;
    int Forward = SUCC(Depot)->Id != Depot->Id + DimensionSaved;
    int Load = Capacity;

    int ColorVisitsCount[COLOR_COUNT] = { 0 };
// an array: How many times each color has been visited
    GainType ColorVisitTimes[COLOR_COUNT] = { 0 };
// an array: The times associated with visiting nodes of each color
    int ReallyArrival[COLOR_COUNT] = { 0 };
    int ColorsVisitedTwice = 0;
// a counter: How many colors have been visited twice

    N = Depot;
    do {
        if (N->Id <= Dim && N != Depot) { // Current node is a customer node
            int colorIndex = N->Color - 1;
```

```
            NextN = Forward ? SUCC(N) : PREDD(N);
            NextN = Forward ? SUCC(NextN) : PREDD(NextN);
            if (N->DraftLimit == 0 && ColorVisitsCount[colorIndex] == 1
&& ColorsVisitedTwice < COLOR_COUNT-1
&& NextN->DepotId == 0){
                // P += Pen1; // no "revisit"
            }
            int physiIndex = N->DraftLimit - 1;
//N->Color = 1 2,..., COLOR_COUNT = #cus

            if (N->ServiceTime == 0 && ColorVisitsCount[colorIndex] < 2){
                ColorVisitsCount[colorIndex] = 2;
                ColorsVisitedTwice ++;
            }
            if (N->ServiceTime != 0 && ColorVisitsCount[colorIndex] == 1){
                ReallyArrival[physiIndex] ++ ;
                ReallyArrival[colorIndex] ++ ;

                ColorVisitsCount[colorIndex] = 2; // Mark the color as visited twice
                ColorsVisitedTwice++; // Increment the count of colors visited twice.
                Load++; // retrieve a drone
                ColorVisitTimes[colorIndex] += N->ServiceTime;
// Add current node's service time to the total visit time for that color.
                if (ColorVisitTimes[colorIndex] > CurrentTime)
// the updated visit time for the color > the current time
                    CurrentTime = ColorVisitTimes[colorIndex];
            }
            if (N->ServiceTime != 0 && ColorVisitsCount[colorIndex] == 0){
                ReallyArrival[physiIndex] ++ ;
                ReallyArrival[colorIndex] ++ ;

                ColorVisitsCount[colorIndex] = 1; // Mark the color as visited once.
                Load--; // launch a drone
                ColorVisitTimes[colorIndex] = CurrentTime + N->ServiceTime;
// Set the visit time for the color
            }
            if (Load > Capacity)
                P += Pen2 * (Load - Capacity);
            if (Load < 0)
                P -= Pen2 * Load;
```

```
            if (P > CurrentPenalty ||
                (P == CurrentPenalty && CurrentGain <= 0)){
                // StartRoute = CurrentRoute;
                return CurrentPenalty + (CurrentGain > 0);
            }

        }
        NextN = Forward ? SUCC(N) : PREDD(N);
        CurrentTime += (C(N, NextN) - N->Pi - NextN->Pi) / Precision;
        N = Forward ? SUCC(NextN) : PREDD(NextN);
    } while (ColorsVisitedTwice < COLOR_COUNT);
    // Go back to the depot
    prevN = Forward ? PREDD(N) : SUCC(N); // current NextN
    prevN2 = Forward ? PREDD(prevN) : SUCC(prevN);
    prevN3 = Forward ? PREDD(prevN2) : SUCC(prevN2); // new
    CurrentTime += (C(Depot, prevN3) - Depot->Pi - prevN3->Pi) / Precision;
    P += CurrentTime;

    P += abs(Load-Capacity)*Pen3;

    for (int i = 0; i < COLOR_COUNT; i++) {
        if (ReallyArrival[i]>2) {
            // P += (ReallyArrival[i]-2)*Pen1; // no "revisit"
        }
    }

    return P;
}
```