Fundamental Performance Limits and Efficient Polices for Transportation-On-Demand Systems

Marco Pavone, Kyle Treleaven, Emilio Frazzoli

Abstract— **Transportation-On-Demand (TOD) systems, where users generate requests for transportation from a pick-up point to a delivery point, are already very popular and are expected to increase in usage dramatically as the inconvenience of privately-owned cars in metropolitan areas becomes excessive. Routing service vehicles through customers is usually accomplished with heuristic algorithms. In this paper we study TOD systems in a formal setting that allows us to characterize fundamental performance limits and devise dynamic routing policies with provable performance guarantees. Specifically, we study TOD systems in the form of a unit-capacity, multiple-vehicle dynamic pick-up and delivery problem, whereby pick-up requests arrive according to a Poisson process and are randomly located according to a general probability density. Corresponding delivery locations are also randomly distributed according to a general probability density, and a number of unit-capacity vehicles must transport demands from their pick-up locations to their delivery locations. We derive insightful fundamental bounds on the steady-state waiting times for the demands, and we devise constant-factor optimal dynamic routing policies. Simulation results are presented and discussed.**

I. Introduction

Transportation-on-demand systems, where users formulate request for transportation from a pick-up point to a delivery point, have become extremely popular. Typical examples are cab-services and dial-a-ride transportation services for the elderly and the disabled. Meanwhile, radically new types of transportation-on-demand systems are being developed, including *Mobility-On-Demand* systems [1], which will provide stacks and racks of light electric vehicles at closely spaced intervals throughout a city: when a person wants to go somewhere, he simply walks to the nearest rack, swipes a card to pick up a vehicle, drives it to the rack nearest to his destination, and drops it off. MOD systems will enable convenient point-to-point travel within urban areas and very high vehicle utilization rates, and will extend availability to those who cannot or do not want to own their own vehicles. Large-scale systems employing traditional, non-electric bicycles have already demonstrated the feasibility of mobility-on-demand in several cities throughout Europe, e.g., Paris, Lyon, Milano, Trento, Zurich and so on [2].

The fundamental problem in transportation-on-demand systems is to route the vehicles with the objective that customers' inconvenience (e.g., in terms of waiting time) is minimized. (In the case of MOD systems, we assume the cars can autonomously drive from a delivery location to the next pick-up location -autonomous driving is an active research topic, see for example [3], [4]-.) This problem falls within the general class of one-to-one Pick-up and Delivery Problems (PDPs), where each customer (or commodity) must be transported from a pick-up site to a delivery site by a fleet of vehicles (with a certain capacity $q \geq 1$). One-to-one PDPs can be either *static* or *dynamic*. In the first case, all requests are known beforehand while in the second case requests are received dynamically and vehicle routes must be adjusted in real-time to meet demand. In some transportation-on-demand systems the setting is static (e.g., for transportation of disabled people the transportation requests are usually formulated a day in advance), however in most scenarios the setting is dynamic. While several exact and heuristic routing algorithms have been studied for static one-to-one PDPs (see [5] for an authoritative survey), few rigorous studies exist for its dynamic counterpart, which often is treated instead by a sequencing of static subproblems. Dynamic one-to-one PDPs can be divided into three main categories [6]: (i) Dynamic Stacker Crane Problem (where the vehicles have unit capacity), (ii) Dynamic Vehicle Routing Problem with Pickups and Deliveries (where the vehicles can transport more than one request), and (iii) the Dynamic Dial-a-Ride Problem (where additional constraints such as time windows are considered). Excellent surveys on heuristics, metaheuristics and online algorithms for Dynamic one-to-one PDPs can be found in [6] and [7]. Even though these algorithms are quite effective in addressing dynamic one-to-one PDPs, alone they do not answer critical questions such as: given a certain number of vehicles, what are the fundamental limits of performance? Is it possible to characterize optimal routing policies? How do customer inconvenience levels scale down as the number of vehicles increases (in other words, what is the marginal benefit of one more vehicle)? How should one pre-position vehicles when there are no outstanding demands?

To the best of our knowledge, the only analytical studies for dynamic one-to-one PDPs are [8] and [9]. Specifically, in [8] the authors consider the *uncapacitated* multiple vehicle case of this problem, and provide lower and upper bounds on the achievable performance. In the same vein, in [9] the authors study the *unit* capacity single vehicle case of this problem, again providing bounds on the achievable performance. The results in [8] and [9] are interesting and insightful, however they are not directly applicable to transportation-on-demand systems, since such systems are characterized by multiple *and* capacitated vehicles.

In this paper we rigorously study routing problems for dynamic transportation-on-demand systems, where pick-up requests arrive according to a Poisson process and are randomly located according to a general probability density. Corresponding delivery locations are also randomly

Marco Pavone, Kyle Treleaven, and Emilio Frazzoli are with the Laboratory for Information and Decision Systems, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, {pavone, ktreleav, frazzoli}@mit.edu.

distributed according to a general probability density, and a *fleet* of *unit-capacity* vehicles must transport demands from their pick-up locations to their delivery locations. The objective is to minimize the expected waiting time for the demands. We assume that the vehicles have single-integrator dynamics and that the environment is a bounded, convex subset within the three-dimensional Euclidean space. These two assumptions are made mainly to ease the exposition: we will in fact argue that the results derived in this paper for this rather artificial but analytically convenient setting hold also for the more realistic setting where vehicles have differential constraints and operate within a two-dimensional manifold (e.g., planar kinematic vehicles with bounded curvature). Our contributions are three-fold: First, we carefully formulate the problem. Second, we establish lower bounds on the expected waiting time in terms of the number of vehicles and other problem's characteristics (e.g., arrival rate of the demands). Finally, we rigorously study a vehicle routing policy whose performance exhibits the same growth rate (in terms of the traffic congestion) as the lower bound.

II. Background material

In this section we summarize the asymptotic properties of the Euclidean traveling salesperson tour and of the bipartite matching problem.

A. The Euclidean Traveling Salesperson Problem

Given a set Q of n points in $\mathbb{R}^d$, the Euclidean traveling salesperson problem (TSP) is to find the minimum-length tour of Q, i.e., the shortest closed path through all points. Let $L_{\text{TSP}}(Q)$ denote the minimum length of a tour through all the points in Q. Assume that the locations of the n points are random variables independently and identically distributed in a compact set Ω according to a density f; in [10] it is shown that there exists a constant β_{TSP} such that, almost surely,

$$\lim_{n\to+\infty} \frac{L_{\text{TSP}}(Q)}{n^{1-1/d}} = \beta_{\text{TSP},d} \int_{\Omega} \bar{f}(q)^{1-1/d}\, dq, \qquad (1)$$

The bound in equation (1) holds for all compact sets Ω, and the shape of Ω only affects the convergence rate to the limit. In [11], the authors note that if Ω is "fairly compact [square] and fairly convex", then equation (1) provides an adequate estimate of the optimal TSP tour length for values of n as low as 15. The constant $\beta_{\text{TSP},3}$ has been estimated numerically as $\beta_{\text{TSP},3} \approx 0.6979 \pm 0.0002$, [12]. Henceforth, we denote $\beta_{\text{TSP},3}$ simply as β_{TSP}.

B. The Bipartite Matching Problem

Let Q be a set of points $X_1, \ldots, X_n$, $Y_1, \ldots, Y_n$ that are i.i.d. in a compact set $\Omega \subset \mathbb{R}^d$, $d \geq 3$, and distributed according to a density f. Let $L_{\text{M}}(Q) = \min_\sigma \sum_{i=1}^n \|X_i - Y_{\sigma(i)}\|$ denote the optimal bipartite matching of the X and Y points, where σ ranges over all permutations of the integers $1, 2, \ldots, n$, and where $\|\cdot\|$ is the Euclidean norm on $\mathbb{R}^d$. In [13] it is shown that there exists a constant β_{M} such that, almost surely,

$$\lim_{n\to+\infty} \frac{L_{\text{M}}(Q)}{n^{1-1/d}} = \beta_{\text{M},d} \int_{\Omega} \bar{f}(q)^{1-1/d}\, dq, \qquad (2)$$

The constant $\beta_{\text{M},3}$ has been estimated numerically as $\beta_{\text{M},3} \approx 0.7080 \pm 0.0002$, [14]. Henceforth, we denote $\beta_{\text{M},3}$ simply as β_{M}.

III. Problem statement

In this section we present a simple yet insightful model for TOD and MOD systems, which takes inspiration from the work on dynamic vehicle routing in [15].

A. The problem

A total of m vehicles operate in a compact, convex environment $\Omega \subset \mathbb{R}^3$. The vehicles are free to move, traveling at a maximum velocity v, within the environment Ω. The vehicles are identical, have unlimited range and are of *unit* capacity (i.e., they can transport one demand at a time). Each vehicle is associated with a depot whose location is $x_k \in \Omega$, $k \in \{1, \ldots, m\}$. Demands are generated according to a homogeneous (i.e., time-invariant) Poisson process, with time intensity $\lambda \in \mathbb{R}_{>0}$. A newly arrived demand has an associated pick-up location which is independent and identically distributed (i.i.d) in Ω according to a density f_{P}. (Note that while a uniform distribution can be a reasonable model for TOD systems, it is not for a MOD system, where pick-ups only happen at specific locations throughout a city.) Each demand must be transported from its pick-up location to its delivery location. The delivery locations are also i.i.d. in Ω according to a density f_{D}. In this paper we will assume that $f_{\text{P}} = f_{\text{D}} = f$. We will also pose the following technical conditions on f [15]:

1) The density f is K-Lipschitz, i.e., $|f(x) - f(y)| \leq K\,\|x - y\|$, $\forall x, y \in \Omega$.
2) The density f is bounded below and above, i.e., $0 < \underline{f} \leq f(x) \leq \overline{f} < \infty$, $\forall x \in \Omega$.

We denote the travel time between the pick-up location of demand j and its delivery location as s_j. A realized demand is removed from the system after one of the vehicles has brought it to its delivery location. Because the sites are generated independently, the expected travel time for demand j is $\mathbb{E}[s_j] = \bar{s} = \frac{1}{v}\int_{x,y\in\Omega} \|y - x\|\, f(x) f(y)\, dy\, dx$. We define the *load factor* $\varrho \doteq \lambda \bar{s}/m$.

The system time of demand j, denoted by T_j, is defined as the elapsed time between the arrival of demand j and the time one of the vehicles completes its service (i.e., it delivers the demand to its delivery location). The waiting time of demand j, W_j, is defined by $W_j = T_j - s_j$. The steady-state system time is defined by $\overline{T} \doteq \limsup_{j\to\infty} \mathbb{E}[T_j]$. A policy for routing the vehicles is said to be *stable* if the expected number of demands in the system is uniformly bounded at all times. A necessary condition for the existence of a stable policy is that $\varrho < 1$; we shall assume $\varrho < 1$ throughout the paper. When we refer to *light-load* conditions, we consider the case $\varrho \to 0^+$, in the sense that $\lambda \to 0^+$; when we refer to *heavy-load* conditions, we consider the case $\varrho \to 1^-$, in the sense that $\lambda \to [m/\bar{s}]^-$.

Let $\mathcal{P}$ be the set of all causal, stable, and stationary policies with the additional (technical) property that decisions occur only at service completion epochs, except for vehicles waiting idle at the depot locations. Letting $\overline{T}_\pi$, denote the system time of a particular policy $\pi \in \mathcal{P}$, the problem is to find a policy π^* (if one exists) such that

$$\overline{T}_{\pi^*} = \inf_{\pi \in \mathcal{P}} \overline{T}_\pi.$$

We let $\overline{T}^*$ denote the infimum of the right hand side above.

We call this problem the *Dynamic Pick-up Delivery problem with m vehicles of unit capacity* (DPDP/m/1).

B. Discussion

A related problem has been previously studied in [9]. In that paper, the DPDP/m/1 is analyzed under the following assumptions: (i) there is only one vehicle (i.e., $m = 1$), (ii) the distribution of pick-up and delivery locations is uniform (i.e., $f = 1/|\Omega|$, where $|\Omega|$ is the area of Ω), and $\Omega \subset \mathbb{R}^2$. First, the authors find a policy that is optimal in light load; then, they derive a lower bound on the system time of the order $(1-\varrho)^{-2}$, and propose a sectoring policy whose bound on the system time is of the order $(1-\varrho)^{-3}$. Finally, they use simulation to analyze other policies. Note that the lower bound is of the order $(1-\varrho)^{-2}$, while the growth rate of the sectoring policy is of the order $(1-\varrho)^{-3}$; therefore, as $\varrho \to 1^-$, the lower bound and the bound for the sectoring policy are arbitrarily far apart.

In the present paper we consider the unit-capacity dynamic Pick-up and Delivery problem in the setting of multiple vehicles with single-integrator dynamics, and arbitrary spatial distribution of demands in three-dimensional environments. Our key contribution is that we are able to find lower bounds and policies that have the *same* growth rate. As mentioned in the introduction, we assume single-integrator dynamics and a three-dimensional environment mostly for analytical convenience: we will argue that the results in this paper hold also for planar vehicles with differential constraints on their motion.

As in many queueing problems, the analysis of the DPDP/m/1 for all values of $\varrho \in (0, 1)$ is difficult. Similarly as in [15], lower bounds for the optimal steady-state system time will be derived for the light-load case (i.e., $\varrho \to 0^+$), and for the heavy-load case (i.e., $\varrho \to 1^-$). Subsequently, policies will be designed for these two limiting regimes, and their performance will be compared to the lower bounds.

We conclude this section by mentioning three major limitations of the DPDP/m/1: (i) the vehicles can freely travel in Ω, i.e., there are no "street constraints", (ii) the delivery locations are *independent* of pick-up locations, and (iii) the densities f_P and f_D are equal.

IV. Lower bounds

In this section we present two lower bounds: the first one is most useful as $\varrho \to 0^+$ (light load), while the second one holds as $\varrho \to 1^-$ (heavy load).

A. A light load lower bound

A lower bound that is most useful in light load (i.e., when $\varrho \to 0^+$) is the following.

Theorem 4.1: The optimal expected time spent in the system by a demand is lower bounded by

$$\overline{T}^* \geq \frac{1}{v}\,\mathbb{E}\left[\min_{i=1}^m \|X - X_i^*\|\right] + \bar{s}. \qquad (3)$$

Proof: The proof is rather straightforward. Assume that we can place the vehicles in the best *a-priori* positions, i.e., at the locations $X_1^*, X_2^*, \dots, X_m^*$, such that $X_1^*, \dots, X_m^* = \arg\min_{X_1,\dots,X_m} \mathbb{E}\left[\min_{i=1}^m \|X - X_i\|\right]$. The expectation is over demand pick-up sites, i.e. X is distributed according to f. We call such a configuration of points an *m-stochastic median*. By definition, the locations $X_1^*, \dots, X_m^*$ minimize the expected distance between the pick-up site of a newly arrived demand and the closest vehicle.

Clearly, the expected time for the vehicle assigned to a newly arrived demand to reach the corresponding pick-up site is at least as large as $\mathbb{E}\left[\min_{i=1}^m \|X - X_i\|\right]/v$. By adding to this the expected time to transfer the demand from its pick-up to its delivery location we obtain the claim. ∎

B. A heavy load lower bound

In this section we present a lower bound that holds as $\varrho \to 1^-$; to derive this bound we make heavy usage of the proof techniques developed in [15]. We start with a definition.

Definition 4.2 (Spatially unbiased policies): Let X be the pick-up location of a randomly chosen demand and W be its wait time. A policy π is said to be *spatially unbiased* if, for every pair of sets $\mathcal{S}_1, \mathcal{S}_2 \subseteq \Omega$, it holds $\mathbb{E}[W\,|X \in \mathcal{S}_1] = \mathbb{E}[W\,|X \in \mathcal{S}_2]$.

In this section we will find a heavy-load lower bound for the class of *unbiased* policies within $\mathcal{P}$.

The expected number of outstanding pick-up sites in an arbitrary region $\mathcal{J}$ of the environment can be expressed as

$$\begin{aligned} N_\mathrm{P}(\mathcal{J}) &= \lambda(\mathcal{J}) W(\mathcal{J}) \\ &= \lambda \int_{\mathcal{J}} f(x)dx\, W = N_\mathrm{P} \int_{\mathcal{J}} f(x)dx, \end{aligned} \qquad (4)$$

where in the first equality we have applied Little's theorem (see [16]) and $W(\mathcal{J}) = W$ because we are considering unbiased policies.

Because of equation (4), and because $f(\cdot)$ is Lipschitz, given a ball $\mathcal{B}(x,z) \doteq \{x' \in \Omega \mid \|x' - x\| \leq z\}$,

$$N_\mathrm{P}(\mathcal{B}(x,z)) = N_\mathrm{P} f(x)\, V_3\, z^3 + N_\mathrm{P}\, o(z^3), \qquad (5)$$

where $V_3 = 3\pi/4$ is the volume of a unit ball in $\mathbb{R}^3$.

In what follows, without loss of generality, we assume that there is a single depot $x_0 \in \Omega$. Let Z be the steady-state expected distance from a vehicle (at the completion epoch of its demand) to the closest outstanding pick-up location, or the depot if closer. We now show a technical lemma, which relates the expected distance $\mathbb{E}[Z]$ to the number of outstanding pick-up locations.

Lemma 4.3: For any unbiased policy in $\mathcal{P}$

$$\lim_{N_\mathrm{P} \to \infty} N_\mathrm{P}^{1/3} \mathbb{E}[Z] \geq \frac{(3/4)^{4/3}}{\sqrt[3]{\pi}} \int_\Omega f^{2/3}(x)\, dx.$$

Proof: We first condition on the event that a randomly tagged demand is delivered at the location $X_\mathrm{D} = x$. Let us fix a neighborhood $\mathcal{D}(N_\mathrm{P}) = \{x' \mid \|x' - x_0\| \leq c^{-1/3}(x)\}$, where $c(x) = N_\mathrm{P}\, V_3\, f(x)$. There are two possible cases.

Case 1: $x \notin \mathcal{D}(N_\mathrm{P})$. For z sufficiently small, i.e., such that $\mathcal{B}(x,z)$ does not contain the depot,

$$\mathbb{P}[Z \leq z|\, X_\mathrm{D} = x] = \mathbb{P}[n_\mathrm{P}^+(\mathcal{B}(x,z)) > 0] \leq N_\mathrm{P}^+(\mathcal{B}(x,z)),$$

where $n_{\rm P}^+$ is the number of outstanding pick-up sites in the ball $\mathcal{B}(x,z)$ at the delivery time of the current demand and $N_{\rm P}^+$ is its expectation. The inequality above holds because $n_{\rm P}^+$ is a non-negative, integer-valued random variable.

For Poisson arrival processes, it holds that $N_{\rm P}^+(\mathcal{J}) = N_{\rm P}(\mathcal{J})$ (this is a consequence of the PASTA property, see [17]), and recalling equation (5) we obtain

$$N_{\rm P}^+(\mathcal{B}(x,z)) = N_{\rm P} f(x)\, V_3\, z^3 + N_{\rm P}\, o(z^3).$$

Hence, we can write

$$\begin{aligned}\mathbb{E}\left[Z \,|\, X_{\rm D}=x\right] &\geq \int_0^{c^{-1/3}(x)} \mathbb{P}[Z>z|\, X_{\rm D}=x] dz\\ &\geq \int_0^{c^{-1/3}(x)} 1 - N_{\rm P} f(x)\, V_3\, z^3 - N_{\rm P}\, o(z^3)\, dz\\ &= \int_0^{c^{-1/3}(x)} 1 - c(x)\, z^3 - N_{\rm P}\, o(z^3)\, dz\\ &= \frac{3}{4}\left[N_{\rm P} V_3\, f(x)\right]^{-1/3} - o(N_{\rm P}^{-1/3}).\end{aligned}$$

Case 2: $x \in \mathcal{D}(N_{\rm P})$. In this case we consider the trivial lower bound $\mathbb{P}[Z > z|\, X_{\rm D} = x] \geq 0$.

We now remove the conditioning on the current delivery site, and we obtain (recall that by assumption f is bounded below by $\underline{f}$, and thus $\int_{\mathcal{D}(N)} dz \leq O(1/N)$)

$$\begin{aligned}&\mathbb{E}\left[Z\right] = \int_\Omega \mathbb{E}\left[Z \,|\, X_{\rm D}=x\right] f(x) dx\\ &\geq \left[N_{\rm P} V_3\right]^{-1/3} \frac{3}{4}\left[\int_{\Omega - \mathcal{D}(N_{\rm P})} f^{-1/3}(x) f(x) dx\right] + \\ &\quad - o(N_{\rm P}^{-1/3})\\ &\geq \left[N_{\rm P} V_3\right]^{-1/3} \frac{3}{4}\left[\int_\Omega f^{2/3}(x) dx - \underline{f}^{-1/3}\overline{f}\int_{\mathcal{D}(N_{\rm P})} dz\right] + \\ &\quad - o(N_{\rm P}^{-1/3})\\ &= \left[N_{\rm P} V_3\right]^{-1/3} \frac{3}{4}\left[\int_\Omega f^{2/3}(x) dx\right] - o(N_{\rm P}^{-1/3}).\end{aligned}$$

Multiplying by $N_{\rm P}^{1/3}$ and taking the limit as $N_{\rm P} \to \infty$, we obtain the claim. ■

We are now in a position to prove the main results of this section.

Theorem 4.4 (Heavy-load lower bound): Within the class of unbiased policies in $\mathcal{P}$

$$\lim_{\varrho\to 1^-} \overline{T}^*(1-\varrho)^3 \geq \gamma_3^3 \frac{\lambda^2}{m^3 v^3}\left[\int_\Omega f^{2/3}(x)\, dx\right]^3$$

where $\gamma_3 \geq (3/4)^{4/3}/\sqrt[3]{\pi}$.

Proof: Let $\mathbb{E}\left[D\right]$ denote the steady-state expected distance traveled empty between the delivery site of a randomly tagged demand and the pick-up site of the next demand to be serviced by the same vehicle. A necessary condition for stability is that

$$\bar{s} + \frac{\mathbb{E}\left[D\right]}{v} \leq \frac{m}{\lambda}. \tag{6}$$

Since, by definition, $\mathbb{E}\left[Z\right] \leq \mathbb{E}\left[D\right]$, equation (6) implies

$$\frac{\lambda}{m}\frac{\mathbb{E}\left[Z\right]}{v} \leq 1-\varrho.$$

By multiplying both sides by $N_{\rm P}^{1/3}$ and raising to the 3-th power we obtain

$$N_{\rm P}(1-\varrho)^3 \geq \frac{\lambda^3 [N_{\rm P}^{1/3}\mathbb{E}\left[Z\right]]^3}{m^3 v^3}.$$

Applying Little's Law, i.e. $N_{\rm P} = \lambda W$, we get

$$\overline{T}(1-\varrho)^3 \geq W(1-\varrho)^3 \geq \frac{\lambda^2 [N^{1/3}\mathbb{E}\left[Z\right]]^3}{m^3 v^3}.$$

Taking the limit as $\varrho \to 1^-$ we trivially have that $N_{\rm P} \to \infty$, hence we can apply lemma 4.3 and obtain the claim. ■

C. Lower bounds with other vehicle's models

It is significant to mention that the order of the lower bounds derived in this section holds for a number of problems in $\mathbb{R}^2$ where service vehicles have more complex dynamics. Consider, for example, Dubins vehicles, which are planar vehicles constrained to move along paths of bounded curvature, without reversing direction and maintaining a constant speed. A Dubins vehicle at position x (with minimum turning radius ρ) has a reachable set $\mathcal{B}_{\text{Dubins}}(x,z)$ with area $z^3/3\rho$ for small distances z, regardless of heading (see [18]). In the heavy load case, to obtain a result similar to theorem 4.4, we simply rewrite equation (5) as $N_{\rm P}\left(\mathcal{B}_{\text{Dubins}}(x,z)\right) = N_{\rm P} f(x)\, z^3/3\rho + N_{\rm P}\, o(z^3)$ and re-apply the analysis of lemma 4.3.

V. Light load policies

A. The m-stochastic median policy

In this section we describe briefly a policy that achieves asymptotic optimality in the light load limit. The policy is the intuitive response to the lower bound construction of (3). For an instance of the problem, we consider the placement of m depots within the environment, at locations $X_1^*, \ldots, X_m^*$, corresponding to the configuration of the m-stochastic median. Each depot will correspond to a queue, and is assigned a service vehicle.

> **The m-stochastic median queue policy (SMQ)** — Upon arrival, a demand is assigned to the depot closest to its pick-up location. The depot's vehicle services its demands in *first-come first-served* order, returning to the depot after each delivery, and waiting there if its queue is empty.

Each of the m resulting queues forms an M/G/1 queue with time intensity $\lambda_i > 0$, such that $\sum_{i=1}^m \lambda_i = \lambda$. By applying the Pollaczek-Khinchin formula for the M/G/1 queue [16], we see that the time spent waiting for the vehicle to service other demands goes to zero as $\lambda_i \to 0^+$, and then

$$\overline{T}_{\text{SMQ},i} \to \frac{1}{v}\mathbb{E}\left[\|X - X_i^*\|\right] + \bar{s}_i,$$

where $\bar{s}_i$ is the expected pickup-to-delivery distance conditioned on the depot. When we remove the conditioning with respect to the depot (note the demand at X serviced by vehicle i implies $\|X - X_i^*\| = \min_{i'=1}^m \|X - X_{i'}^*\|$), and take $\lambda \to 0^+$, we find that the expected waiting time under this policy approaches exactly

$$\overline{T}_{\text{SMQ}} \to \frac{1}{v}\mathbb{E}\left[\min_{i=1}^m \|X - X_i^*\|\right] + \bar{s},$$

showing the tightness of the lower bound.

VI. Heavy load policies

Before presenting and analyzing a policy that is particularly effective in heavy load, we define the concept of the *bipartite matching tour*.

A. Bipartite matching tour

Let $X_0, X_1, \ldots, X_n, Y_1, \ldots, Y_n$ be points in $\mathbb{R}^d$. The point X_0 will represent the initial location of a vehicle, the points $\mathcal{X} = \{X_1, \ldots, X_n\}$ will be pick-up locations and the points $\mathcal{Y} = \{Y_1, \ldots, Y_n\}$ will be the corresponding delivery locations. A bipartite matching tour is essentially an approximation of a shortest length tour through the points $X_0, \mathcal{X}, \mathcal{Y}$ with the constraint that when a vehicle visits a pick-up point, the next point to be visited is the corresponding delivery point (such a tour is know in the literature as the *stacker crane tour*).

The bipartite matching tour is constructed as follows. First we add n directed edges $X_i \to Y_i$ that connect pick-up locations to the corresponding delivery locations. Second, we find a bipartite matching for the $\mathcal{X}$ and $\mathcal{Y}$ locations. By adding the n edges of the bipartite matching to the n pick-up to delivery edges $X_i \to Y_i$ we obtain one or more tours, which we call *secondary tours*. Finally, we find a TSP tour (which we call the *primary tour*) across the locations $X_0, X_1, \ldots, X_n$ and we add the corresponding edges. A bipartite matching tour is then as follows: we start at X_0, and follow the primary tour until the first location in $\mathcal{X}$ is reached, say X_j. Then, we follow the secondary tour starting at X_j until we reach again X_j. We resume the primary tour and follow it until we find the next unvisited point in $\mathcal{X}$, say X_k. The procedure is iterated in this way until we reach X_0 again (see figure 2). This concept was originally introduced in [19].

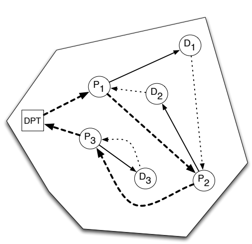

Fig. 1. A bipartite matching tour. The square represents the current location of the vehicle. P_1, P_2, P_3 are pick-up locations and D_1, D_2, D_3 are the corresponding delivery locations. Solid arrows show links between pick-up and delivery sites. Dotted arrows show links obtained by the bipartite matching between delivery and pick-up sites. Finally, dashed arrows show the primary tour (TSP) through pick-up sites. The bipartite matching tour is: $DPT \to P_1 \to D_1 \to P_2 \to D_2 \to P_1 \to P_2 \to P_3 \to D_3 \to P_3 \to DPT$.

B. The randomized batch policy

In this section we present an *unbiased* service policy, which we call randomized batch policy (RB).

> **The randomized batch policy (RB) —** Each newly arrived demand is assigned with probability $p_k = 1/m$ to the vehicle k, $k \in \{1, \ldots, m\}$. Then, for each vehicle k: Let D_k be the set of outstanding demands waiting for service. If $D_k = \emptyset$, move to x_k (the depot). If, instead, $D_k \neq \emptyset$, compute a bipartite matching tour through the current vehicle position and all demands in D_k and service all demands by following such tour. Repeat.

C. Analysis

The performance of the RB policy in heavy load is characterized by the following theorem.

Theorem 6.1 (Performance of RB policy in heavy load): As $\varrho \to 1^-$, the system time for the RB policy satisfies

$$\overline{T} \leq \frac{\lambda^2\left(\beta_{\text{TSP}} + \beta_{\text{M}}\right)^3 \left(\int_\Omega f^{2/3}(x)\,dx\right)^3}{v^3\, m^2\,(1-\varrho)^3} \tag{7}$$

The proof of Theorem 6.1 builds on a number of intermediate results; we start with the following lemma, similar to Lemma 1 in [20], characterizing the number of outstanding demands in heavy load.

Lemma 6.2 (Number of demands in heavy load): In heavy load (i.e., $\varrho \to 1^-$), after a transient, the number of demands serviced in a single tour is very large with high probability (i.e., the number of demands tends to $+\infty$ with probability that tends to 1, as ϱ approaches 1^-).

Proof: The proof is very similar to the one of lemma 4.3 in [21] and thus it is omitted. ■

In this policy, each vehicle sees a demand arrival process which is Poisson with rate λ/m and operates within the entire workspace Ω. Thus, the m-vehicle problem has been turned into m independent and (statistically) identical single-vehicle problems, each with a Poisson arrival process with rate λ/m. As a consequence, we have $\mathbb{E}\left[T \,|\, \text{demand assigned to vehicle } j\right] = \mathbb{E}\left[T \,|\, \text{demand assigned to vehicle } k\right]$ and

$$\begin{aligned}\overline{T}_{\text{RB}} &= \sum_{k=1}^m \frac{1}{m}\,\mathbb{E}\left[T \,|\, \text{demand assigned to vehicle } k\right] \\ &= \mathbb{E}\left[T \,|\, \text{demand assigned to vehicle } 1\right].\end{aligned} \tag{8}$$

Therefore it is enough to study the system time for the demands assigned to vehicle 1. For simplicity of notation we omit the label 1 in what follows.

Lemma 6.2 has two implications. First, since the number of demands is very large at the time instants when the vehicle starts a new bipartite matching tour, we can apply equation (1) to estimate the length of the TSP tour and equation (2) to estimate the length of the bipartite matching. Second, since $D \neq \emptyset$ with high probability, the policy operates with the condition $D = \emptyset$ *always* false.

We refer to the time instant t_i, $i \geq 0$, in which the vehicle starts a new bipartite matching tour as the epoch i of the policy; we refer to the time interval between epoch i and epoch $i+1$ as the ith iteration. Let n_i be the number of outstanding demands serviced during iteration i. Finally, let C_i be the time interval between the time instant the vehicle starts to service demands during iteration i and the time instant the vehicle starts to service demands during next iteration $i+1$. Demands arrive according to a Poisson process with rate $\hat{\lambda} \doteq \lambda/m$; then, we have $\mathbb{E}[n_{i+1}] = \hat{\lambda}\,\mathbb{E}[C_i]$. The time interval C_i is equal to the time to traverse the bipartite matching tour through the n_i demands, which in turn is the sum of three components:

1) the time to traverse the edges of the TSP tour;
2) the time to traverse the edges of the bipartite matching;
3) the n_i travel times from pick-up locations to delivery locations.

Assume that i is large enough (say, $i \geq \bar{i}$) so that, according to Lemma 6.2, the number of outstanding demands is large. Then, the expected time to traverse the bipartite matching tour through the n_i demands, which we call $T(n_i)$, can be upper bounded as

$$\begin{aligned}\mathbb{E}[T(n_i)] &= \frac{1}{v}\mathbb{E}[L_{\mathrm{TSP}}(n_i+1)] + \frac{1}{v}\mathbb{E}[L_{\mathrm{M}}(n_i)] + \mathbb{E}\left[\textstyle\sum_{j=1}^{n_i} s_j\right] \\ &\leq \mathbb{E}[n_i]^{2/3}\,\frac{\beta_{\mathrm{TSP}}+\beta_{\mathrm{M}}}{v}\int_\Omega f^{2/3}(x)\,dx + \mathbb{E}[n_i]\,\bar{s} + O(1),\end{aligned} \tag{9}$$

where we use equation (1) and we apply Jensen's inequality for concave functions in the form $\mathbb{E}\left[X^{2/3}\right] \leq \mathbb{E}[X]^{2/3}$.

Then, we obtain the following recurrence relation (where we define $\bar{n}_i \doteq \mathbb{E}[n_i]$):

$$\begin{aligned}\bar{n}_{i+1} &= \hat{\lambda}\mathbb{E}[C_i] \\ &\leq \hat{\lambda}\Big(\bar{n}_i^{2/3}\frac{\beta_{\mathrm{TSP}}+\beta_{\mathrm{M}}}{v}\int_\Omega f^{2/3}(x)\,dx + \bar{n}_i\bar{s} + O(1)\Big).\end{aligned} \tag{10}$$

The above inequality describes a system of recurrence relations that allows us to find an upper bound on $\limsup_{i\to+\infty}\bar{n}_i$. The following lemma bounds the value to which the limit $\limsup_{i\to+\infty}\bar{n}_i$ converges.

Lemma 6.3 (Steady state number of demands): In heavy load, for every initial condition $\bar{n}_1$, the trajectory $i \mapsto \bar{n}_i$ satisfies

$$\bar{n} \doteq \limsup_{i\to+\infty}\bar{n}_i \leq \frac{\lambda^3(\beta_{\mathrm{TSP}}+\beta_{\mathrm{M}})^3\Big(\int_\Omega f^{2/3}(x)\,dx\Big)^3}{v^3\,m^3\,(1-\varrho)^3}.$$

Proof: By Lemma 6.2, n_i tends to $+\infty$ with probability that tends to 1, as ϱ approaches 1^-; then, after a transient, the term $O(1)$ is negligible compared to the other terms in the right hand side of equation (10), and therefore it can be ignored (its inclusion in the proof is conceptually straightforward, but makes the analysis more cumbersome).

Next we define two auxiliary systems, System-Y and System-Z. We define System-Y (with state $y \in \mathbb{R}$) as

$$y(i+1) = \hat{\lambda}\Big(y(i)^{2/3}\frac{\beta_{\mathrm{TSP}}+\beta_{\mathrm{M}}}{v}\int_\Omega f^{2/3}(x)\,dx + y(i)\bar{s}\Big). \tag{11}$$

System-Y is obtained by replacing the inequality in equation (10) with an equality. Pick, now, any $\varepsilon > 0$. From Young's inequality

$$a = \frac{a\,(p\,\varepsilon)^\alpha}{(p\,\varepsilon)^\alpha} \leq \Big(a\,(p\,\varepsilon)^\alpha\Big)^p\frac{1}{p} + \Big(\frac{1}{(p\,\varepsilon)^\alpha}\Big)^q\frac{1}{q},$$

where $a \in \mathbb{R}_{\geq 0}$, $p, q \in \mathbb{R}_{>0}$, $1/p + 1/q = 1$ and $\alpha, \varepsilon \in \mathbb{R}_{>0}$. By letting $a = y^{2/3}$, $p = 3/2$, $q = 3$ and $\alpha = 2/3$ we obtain:

$$y^{2/3} \leq \varepsilon\,y + \frac{4}{27\,\varepsilon^2}.$$

By applying the above inequality in equation (11) we obtain

$$\begin{aligned}y(i+1) \leq &\hat{\lambda}\Big(\bar{s} + \varepsilon\,\frac{\beta_{\mathrm{TSP}}+\beta_{\mathrm{M}}}{v}\int_\Omega f^{2/3}(x)\,dx\Big)\,y(i) \\ &+ \underbrace{\frac{4\,\hat{\lambda}}{27\,\varepsilon^2}\frac{\beta_{\mathrm{TSP}}+\beta_{\mathrm{M}}}{v}\int_\Omega f^{2/3}(x)\,dx}_{=O(1)\,\hat{\lambda}/\varepsilon^2}.\end{aligned} \tag{12}$$

Next, define System-Z as

$$\begin{aligned}z(i+1) = &\hat{\lambda}\Big(\bar{s} + \varepsilon\,\frac{\beta_{\mathrm{TSP}}+\beta_{\mathrm{M}}}{v}\int_\Omega f^{2/3}(x)\,dx\Big)\,z(i) \\ &+ O(1)\,\frac{\hat{\lambda}}{\varepsilon^2}.\end{aligned} \tag{13}$$

It is immediate to show that if $\bar{n}_{\bar{i}} \leq y(\bar{i}) \leq z(\bar{i})$, then

$$\bar{n}_i \leq y(i) \leq z(i), \quad \text{for all } i \geq \bar{i}. \tag{14}$$

(Note that System-Y and System-Z are virtual systems for which we can arbitrarily pick the initial conditions.) The proof now proceeds as follows. First, we show that the trajectories of System-Z are bounded; this fact, together with equation (14), implies that also trajectories of variables $\bar{n}_i$ and $y(i)$ are bounded. Then, we will compute $\limsup_{i\to+\infty} y(i)$; this quantity, together with equation (14), will yield the desired result.

Let us consider the first issue, namely boundedness of trajectories for System-Z (described in equation (13)). System-Z is a discrete-time linear system and can be rewritten in compact form as

$$z(i+1) = \Big(\varrho + \varepsilon\,b\Big)z(i) + O(1)\,\frac{\hat{\lambda}}{\varepsilon^2},$$

where $\varrho = \hat{\lambda}\bar{s}$ and $b = \hat{\lambda}(\beta_{\mathrm{TSP}}+\beta_{\mathrm{M}})\int_\Omega f^{2/3}(x)\,dx/v$. Since, by assumption, $\varrho < 1$, there exists a sufficiently small $\varepsilon > 0$ such that $\varrho + \varepsilon\,b < 1$. Accordingly, having selected a sufficiently small ε, each solution $i \mapsto z(i) \in \mathbb{R}_{\geq 0}$ of System-Z converges exponentially fast to the unique equilibrium point

$$z^*(\varepsilon) = \Big(1 - \varrho - \varepsilon\,b\Big)^{-1}O(1)\,\frac{\hat{\lambda}}{\varepsilon^2}. \tag{15}$$

Combining equation (14) with the previous statement, we see that the solutions $i \mapsto \bar{n}_i$ and $i \mapsto y(i)$ are bounded. Thus

$$\limsup_{i\to+\infty}\bar{n}_i \leq \limsup_{i\to+\infty} y(i) < +\infty. \tag{16}$$

Finally, we turn our attention to the computation of $y \doteq \limsup_{i\to+\infty} y(i)$. Taking the $\limsup$ of the left- and right-hand sides of equation (11), and noting that $\limsup_{i\to+\infty} y^{2/3}(i) = \Big(\limsup_{i\to+\infty} y(i)\Big)^{2/3}$, since $x \to x^{2/3}$ is continuous and strictly monotone increasing on $\mathbb{R}_{>0}$, we obtain that

$$y = y^{2/3}\,\hat{\lambda}\,\frac{\beta_{\mathrm{TSP}}+\beta_{\mathrm{M}}}{v}\int_\Omega f^{2/3}(x)\,dx + y\,\varrho; \tag{17}$$

rearranging we obtain

$$y = \frac{\hat{\lambda}^3(\beta_{\rm TSP}+\beta_{\rm M})^3\Big(\int_\Omega f^{2/3}(x)\,dx\Big)^3}{v^3\,(1-\varrho)^3}.$$

Noting that from equation (16) $\limsup_{i\to+\infty}\bar{n}_i \le y$, we obtain the desired result. ∎

We are now in a position to prove Theorem 6.1.

Proof: [Proof of Theorem 6.1] Define $\overline{C} \doteq \limsup_{i\to\infty}\mathbb{E}\,[C_i]$; then we have, by using the upper bound on $\mathbb{E}\,[C_i]$ in equation (10) (neglecting $O(1)$ terms),

$$\overline{C} \doteq \limsup_{i\to\infty}\mathbb{E}\,[C_i] \le \Big(\bar{n}^{2/3}\frac{\beta_{\rm TSP}+\beta_{\rm M}}{v}\int_\Omega f^{2/3}(x)\,dx + \bar{n}\,\bar{s}\Big) \le \frac{\lambda^2(\beta_{\rm TSP}+\beta_{\rm M})^3\Big(\int_\Omega f^{2/3}(x)\,dx\Big)^3}{v^3\,m^2\,(1-\varrho)^2} + \frac{\lambda^2\,\varrho\,(\beta_{\rm TSP}+\beta_{\rm M})^3\Big(\int_\Omega f^{2/3}(x)\,dx\Big)^3}{v^3\,m^2\,(1-\varrho)^3}.$$

Hence, in the limit $\varrho\to 1^-$, we have $\overline{C} \le \lambda^2\,(\beta_{\rm TSP}+\beta_{\rm M})^3\Big(\int_\Omega f^{2/3}(x)\,dx\Big)^3\Big/ v^3\,m^2\,(1-\varrho)^3$.

The expected steady-state system time of a random demand, $\overline{T}$, is then upper bounded, as $\varrho\to 1^-$, by

$$\begin{aligned}\overline{T} &\le \frac{1}{2}\,\overline{C} + \frac{1}{2}\bar{s}\,\bar{n}\\ &\le \frac{1}{2}\,\overline{C} + \frac{1}{2}\frac{\lambda^2\,(\beta_{\rm TSP}+\beta_{\rm M})^3\Big(\int_\Omega f^{2/3}(x)\,dx\Big)^3}{v^3\,m^2\,(1-\varrho)^3},\end{aligned}\tag{18}$$

where we used the fact that, as $\varrho\to 1^-$, the travel time along the bipartite matching tour is negligible compared to the pick-up to delivery transfer times. Collecting the results we obtain the claim. ∎

D. RB policy for Dubins vehicles in $\mathbb{R}^2$

The RB policy, with some modifications, can be adapted to handle the case where vehicles have more complex dynamics. Consider, for example, Dubins vehicles in $\mathbb{R}^2$. We can adapt the concept of bipartite matching tour as follows. Assume there are n outstanding demands. First, we find a Dubins TSP tour (i.e. a TSP tour that respects the differential constraint of a Dubins vehicle) through the pick-up sites and the vehicle's current location. In [22], an algorithm is proposed to construct such a tour, based on a tiling of the plane into $2n$ "beads", geometric shapes adapted to the Dubins dynamics (see [22] for a rigorous description of a bead tiling). The algorithm returns a tour whose length is of order $n^{2/3}$. Note that the Dubins TSP induces a heading constraint for the pick-up sites. Second, we consider again a bead tiling of the plane into $2n$ beads, and we find a minimum length, maximum cardinality bipartite matching from delivery sites to bead entrance points. The assignment induces heading constraints for delivery sites, and by combining theorems 4.1 and 4.8 in [22] the length of this matching is of order $n^{2/3}$. We then find a minimum length, maximum cardinality bipartite matching from bead entrance points to pick-up sites. By combining again theorems 4.1 and 4.8 in [22], the length of this matching is of order $n^{2/3}$. At this stage, each entrance point is associated with one delivery point and with one pick-up point, hence we have found a bipartite matching from delivery points to pick-up points whose length is of order $n^{2/3}$. Finally, we find minimum length paths from the pick-up sites to their corresponding delivery sites with constrained heading. Since the sum of the lengths of the Dubins TSP and of the (Dubins) bipartite matching is of the order $n^{2/3}$, the analysis in the proof of theorem 6.1 holds, and a theorem analogous to theorem 6.1 can be stated for Dubins vehicles in $\mathbb{R}^2$.

E. Comparison with the lower bound

With theorem 6.1 we can readily prove that the steady-state system time under the RB policy differs from the heavy-load lower bound in theorem 4.4 by a known *constant* factor; specifically, the system time under the RB policy has the *same* growth rate of the lower bound (this is not the case in the work [9]).

Theorem 6.4: Let $\overline{T}^*$ be the optimal system time within the class of unbiased policies in $\mathcal{P}$; then

$$\frac{\overline{T}_{\rm RB}}{\overline{T}^*} \le m\,\frac{(\beta_{\rm TSP}+\beta_{\rm M})^3}{\gamma_3^3}, \quad \text{as } \varrho\to 1^-.$$

VII. Simulation

In this section we present simulation results for the RB policy. We mention that we also studied by simulation several other unbiased policies (e.g., policies for which the demand assignment is not random but follows more advance criteria), however the performance of these policies was similar to that of the RB policy, and so they will not be discussed.

In all simulations we assumed the environment Ω to be the unit cube $[0,\,1]^3$ and the spatial demand density f to be uniform over Ω. For each set of parameters (e.g., ϱ, m etc.) we generated 20 instances by simulation, and computed the mean demand system time.

Simulations of the RB policy were performed using `linkern`[1] as a solver to generate approximations to the optimal TSP tour. A Python implementation of the Kuhn-Munkres assignment algorithm [23] was used to generate Euclidean bipartite matchings.

In figure 2(a) we show the dependance of the system time $\overline{T}_{\rm RB}$ on the load factor ϱ with a number of vehicles $m=3$. We consider values of $\varrho\in[0.6,\,0.75]$, which correspond to a moderate/heavy load. One can observe that the experimental results are within the theoretical lower and upper bounds (even though these bounds formally hold only in the limit $\varrho\to 1^-$); moreover, one can observe that the performance of the RB policy is significantly better than what is predicted by the upper bound; hence we believe that the upper bound in theorem 6.1 is rather conservative. We also study how $\overline{T}_{\rm RB}$ scales with m. To this end, we set the load factor $\varrho=0.6$ and we simulate the RB policy with $m=1,2,3,4,5$. Note that by fixing ρ, we are implicitly letting λ be a function of m, since, by definition, $\lambda=\varrho\,m/\bar{s}$. Hence, by increasing m while keeping ϱ fixed the upper bound in theorem 6.1 stays constant, while the lower bound in theorem 4.4 scales

[1]The TSP solver `linkern` is freely available for academic research use at `http://www.tsp.gatech.edu/concorde.html`.

as $1/m$. Figure 2(b) reports the value of $\overline{T}_{\mathrm{RB}}$ for each value of m and show that $\overline{T}_{\mathrm{RB}}$ stays constant. Hence, recalling our previous discussion, we argue that the RB policy indeed scales as $1/m^2$.

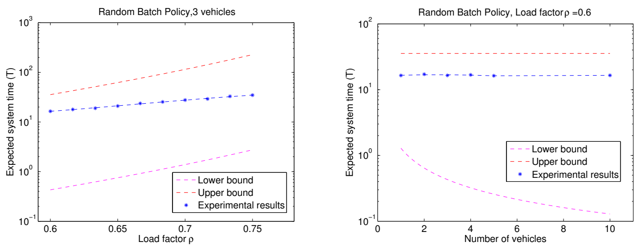

Fig. 2. Performance of RB policy and comparison with upper and lower bounds. Left figure: $\overline{T}_{\mathrm{RB}}$ versus ϱ. Right figure: scaling of $\overline{T}_{\mathrm{RB}}$ with respect to m.

VIII. Conclusion

In this paper we studied a dynamic PDP with multiple vehicles of unit capacity and we argued that this is a reasonable model for TOD and MOD systems. We presented a policy that is optimal in light load and we showed that in heavy load the system time under the RB policy is, asymptotically, within a constant factor of the optimal performance. An open issue in this context is that while the lower bound scales with the number of vehicles as $O(1/m^3)$, the upper bound on the RB policy scales as $O(1/m^2)$. Hence, the optimal scaling of the system time with respect to the number of vehicles is between $O(1/m^3)$ and $O(1/m^2)$, but the exact value is still unknown.

This paper leaves numerous important extensions open for further research. First, it is of strong economic interest to precisely characterize the optimal scaling of the system time with respect to the number of vehicles. This goal would require a tighter lower bound (we conjecture, however, that the scaling $O(1/m^3)$ is indeed correct) and/or devising a policy with a better scaling in terms of m. Second, our initial motivation was to study TOD and MOD systems, for which f_P and f_D might indeed be drastically different (even with different support). Hence we plan to extend our analysis to the case $f_P \neq f_D$. Finally, we plan to consider impatient demands that disappear if they are not serviced within a certain time window.

Acknowledgments

This research was supported in part by the Future Urban Mobility project of the Singapore-MIT Alliance for Research and Technology (SMART) Center, with funding from Singapore's National Research Foundation. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the supporting organizations.

References

[1] W. J. Mitchell, C. E. Borroni-Bird, and L. D. Burns. *Reinventing the Automobile*. MIT Press, 2010.

[2] P. Midgley. The role of smart bike-sharing systems in urban mobility. Available at http://www.gtkp.org/uploads/public/documents/Knowledge/The-Role-of-Smart-Bike.pdf, 2009.

[3] M. Buehler, K. Iagnemma, and S. Singh, editors. *The DARPA Urban Challenge: Autonomous Vehicles in City Traffic*. Number 56 in Springer Tracts in Advanced Robotics. Springer, 2009.

[4] M. Buehler, K. Iagnemma, and S. Singh, editors. *The 2005 DARPA Grand Challenge: the Great Robot Race*. Number 36 in Springer Tracts in Advanced Robotics. Springer, 2007.

[5] J. F. Cordeau, G. Laporte, J. Y. Potvin, and M.W.P. Savelsbergh. Transportation on demand. In C. Barnhart and G. Laporte, editors, *Transportation, Handbooks in Operations Research and Management Science*, volume 14, pages 429–466. Elsevier, Amsterdam, The Netherlands, 2007.

[6] G. Berbeglia, J. F. Cordeau, and G. Laporte. Dynamic pickup and delivery problems. *European Journal of Operational Research*, 202(1):8 – 15, 2010.

[7] S. N. Parragh, K. F. Doerner, and R. F. Hartl. A survey on pickup and delivery problems. *Journal fur Betriebswirtschaft*, 58(2):81 – 117, 2008.

[8] H.A. Waisanen, D. Shah, and M.A. Dahleh. A dynamic pickup and delivery problem in mobile networks under information constraints. *IEEE Transactions on Automatic Control*, 53(6):1419–1433, 2008.

[9] M. R. Swihart and J. D. Papastavrou. A stochastic and dynamic model for the single-vehicle pick-up and delivery problem. *European Journal of Operational Research*, 114(3):447–464, 1999.

[10] J. M. Steele. Probabilistic and worst case analyses of classical problems of combinatorial optimization in Euclidean space. *Mathematics of Operations Research*, 15(4):749, 1990.

[11] R. C. Larson and A. R. Odoni. *Urban Operations Research*. 1981.

[12] A. G. Percus and O. C. Martin. Finite size and dimensional dependence of the Euclidean traveling salesman problem. *Physical Review Letters*, 76(8):1188–1191, 1996.

[13] V. Dobric and J. E. Yukich. Asymptotics for transportation cost in high dimensions. *Journal of Theoretical Probability*, 8(1):97–118, 1995.

[14] J. Houdayer, J.H. Boutet de Monvel, and O.C. Martin. Comparing mean field and euclidean matching problems. *The European Physical Journal B - Condensed Matter and Complex Systems*, 6(3):383–393, 1998.

[15] D. J. Bertsimas and G. J. van Ryzin. Stochastic and dynamic vehicle routing with general interarrival and service time distributions. *Advances in Applied Probability*, 25:947–978, 1993.

[16] R. G. Gallager. *Discrete Stochastic Processes*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1996.

[17] R. W. Wolff. Poisson arrivals see time averages. *Operations Research*, 30(2):223–231, 1982.

[18] J. Enright, K. Savla, and E. Frazzoli. Coverage control for nonholonomic agents. In *47th IEEE Conference on Decision and Control, 2008. CDC 2008*, pages 4250–4256, 2008.

[19] H. Xu. *Optimal policies for stochastic and dynamic vehicle routing problems.* Dept. of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA, 1995.

[20] J. D. Papastavrou. A stochastic and dynamic routing policy using branching processes with state depended immigration. *European Journal of Operational Research*, 95:167–177, 1996.

[21] M. Pavone, E. Frazzoli, and F. Bullo. Distributed and adaptive algorithms for vehicle routing in a stochastic and dynamic environment. *IEEE Trans. on Automatic Control*, 2009. Provisionally Accepted. Available at `http://arxiv.org/abs/0903.3624`.

[22] K. Savla, E. Frazzoli, and F. Bullo. Traveling salesperson problems for the dubins vehicle. *IEEE Transactions on Automatic Control*, 53(6):1378 –1391, july 2008.

[23] J. Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38, March 1957.