

MIT Open Access Articles

Fundamental performance limits and efficient polices for Transportation-On-Demand systems

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Pavone, Marco, Kyle Treleaven, and Emilio Frazzoli. "Fundamental Performance Limits and Efficient Polices for Transportation-On-Demand Systems." 49th IEEE Conference on Decision and Control (CDC). Atlanta, GA, USA, 2010. 5622-5629. © 2011 IEEE

As Published: <http://dx.doi.org/10.1109/CDC.2010.5717552>

Publisher: Institute of Electrical and Electronics Engineers

Persistent URL: <http://hdl.handle.net/1721.1/65388>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Fundamental Performance Limits and Efficient Policies 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 requests 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

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, MA 02139 {pavone, ktreleav, frazzoli}@mit.edu.

randomly located according to a general probability density. Corresponding delivery locations are also randomly 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. 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 propose two vehicle routing policies and analyze their performance: one of the policies is optimal in light load (i.e., when the arrival rate for the demands is small), while the other one performs within a constant factor of the lower bound in heavy load (i.e., when the arrival rate for the demands is large), where the constant is independent of the number of vehicles, the arrival rate of demands, and the spatial density for the demand locations.

II. BACKGROUND MATERIAL

In this section we summarize the asymptotic properties of the Euclidean traveling salesperson tour and of the bipartite matching problem; these properties will be fundamental to analyze the heavy-load policy we study in section VI.

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 (i.i.d.) in a compact set $\Omega \subset \mathbb{R}^d$ according to a density f ; in [10] it is shown that there exists a constant β_{TSP} such that, almost surely,

$$\lim_{n \rightarrow +\infty} \frac{L_{\text{TSP}}(Q)}{n^{1-1/d}} = \beta_{\text{TSP},d} \int_{\Omega} \bar{f}(x)^{1-1/d} dx. \quad (1)$$

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 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 $X_1, \dots, X_n, Y_1, \dots, Y_n$ 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 $L_M(Q) = \min_{\sigma} \sum_{i=1}^n \|X_i - Y_{\sigma(i)}\|$ denote the optimal bipartite matching of the X and Y points, where σ ranges over all permutations of the integers $1, 2, \dots, n$, 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,

$$\lim_{n \rightarrow +\infty} \frac{L_M(Q)}{n^{1-1/d}} = \beta_{M,d} \int_{\Omega} \bar{f}(x)^{1-1/d} dx. \quad (2)$$

The constant $\beta_{M,3}$ has been estimated numerically as $\beta_{M,3} \approx 0.7080 \pm 0.0002$, [14]. Henceforth, we denote $\beta_{M,3}$ 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 $x_k \in \Omega$, $k \in \{1, \dots, m\}$. 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 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 $f_P = f_D = f$. 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., $0 < \underline{f} \leq f(x) \leq \bar{f} < \infty$, $\forall x \in \Omega$.

A realized demand is removed from the system after one of the vehicles has transported it to its delivery location.

We denote the travel time between the pick-up location of demand j and its delivery location as s_j (demands are labeled in an increasing order with respect to time of arrival). Because the sites are generated independently, the expected travel time for demand j is $\bar{s} \doteq \mathbb{E}[s_j] = \frac{1}{v} \int_{x,y \in \Omega} \|y - x\| f(x)f(y) dy dx$. We define the *load factor* $\rho \doteq \lambda \bar{s} / m$.

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 $\bar{T} \doteq \limsup_{j \rightarrow \infty} \mathbb{E}[T_j]$. 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 $\rho < 1$; we shall assume $\rho < 1$ throughout the paper. When we refer to *light-load* conditions, we consider

the case $\varrho \rightarrow 0^+$, in the sense that $\lambda \rightarrow 0^+$; when we refer to *heavy-load* conditions, we consider the case $\varrho \rightarrow 1^-$, in the sense that $\lambda \rightarrow [m/\bar{s}]^-$.

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 \bar{T}_π 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

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

We let \bar{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 (iii) $\Omega \subset \mathbb{R}^2$. 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 $\varrho \rightarrow 1^-$, 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 exhibit the *same* growth rate in terms of all problem's characteristics (e.g., arrival rate of demands and number of vehicles).

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., $\varrho \rightarrow 0^+$), and for the heavy-load case (i.e., $\varrho \rightarrow 1^-$). 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 $\varrho \rightarrow 0^+$ (light load), while the second one holds as $\varrho \rightarrow 1^-$ (heavy load).

A. A light load lower bound

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

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

$$\bar{T}^* \geq \frac{1}{v} \min_{(x_1, \dots, x_m) \in \Omega^m} \mathbb{E} [\min_{i=1}^m \|X - x_i\|] + \bar{s}, \quad (3)$$

where the expectation is with respect to the density f .

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^*, \dots, x_m^* , such that $x_1^*, \dots, x_m^* = \arg \min_{x_1, \dots, x_m} \mathbb{E} [\min_{i=1}^m \|X - x_i\|]$. 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 $\mathbb{E} [\min_{i=1}^m \|X - x_i^*\|]/v$. 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 $\varrho \rightarrow 1^-$; 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{S} of the environment can be expressed as

$$\begin{aligned} \bar{N}_P(\mathcal{S}) &= \lambda(\mathcal{S}) \bar{W}(\mathcal{S}) \\ &= \lambda \int_{\mathcal{S}} f(x) dx \bar{W} = \bar{N}_P \int_{\mathcal{S}} f(x) dx, \end{aligned} \quad (4)$$

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

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

$$\bar{N}_P(\mathcal{B}(x, z)) = \bar{N}_P f(x) V_3 z^3 + \bar{N}_P o(z^3), \quad (5)$$

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

In what follows, to ease the exposition, we assume that there is a single depot $x_0 \in \Omega$ (extension to the general case is straightforward but cumbersome). Let Z be the steady-state distance from a vehicle (at the completion epoch of its last assigned 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}

$$\lim_{\bar{N}_P \rightarrow \infty} \bar{N}_P^{1/3} \mathbb{E}[Z] \geq \frac{(3/4)^{4/3}}{\sqrt[3]{\pi}} \int_{\Omega} f^{2/3}(x) dx.$$

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

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

$$\mathbb{P}[Z \leq z \mid X_D = x] = \mathbb{P}[N_P^+(\mathcal{B}(x, z)) > 0] \leq \bar{N}_P^+(\mathcal{B}(x, z)),$$

where N_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 \bar{N}_P^+ is its expectation. The inequality above holds because N_P^+ is a non-negative, integer-valued random variable.

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

$$\bar{N}_P^+(\mathcal{B}(x, z)) = \bar{N}_P f(x) V_3 z^3 + \bar{N}_P o(z^3).$$

Hence, we can write

$$\begin{aligned} \mathbb{E}[Z \mid X_D = x] &\geq \int_0^{c^{-1/3}(x)} \mathbb{P}[Z > z \mid X_D = x] dz \\ &\geq \int_0^{c^{-1/3}(x)} 1 - \bar{N}_P f(x) V_3 z^3 - \bar{N}_P o(z^3) dz \\ &= \int_0^{c^{-1/3}(x)} 1 - c(x) z^3 - \bar{N}_P o(z^3) dz \\ &= \frac{3}{4} [\bar{N}_P V_3 f(x)]^{-1/3} - o(\bar{N}_P^{-1/3}). \end{aligned}$$

Case 2: $x \in \mathcal{D}(\bar{N}_P)$. In this case we consider the trivial lower bound $\mathbb{P}[Z > z \mid X_D = x] \geq 0$.

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 $\int_{\mathcal{D}(\bar{N}_P)} dz \leq O(1/\bar{N}_P)$)

$$\begin{aligned} \mathbb{E}[Z] &= \int_{\Omega} \mathbb{E}[Z \mid X_D = x] f(x) dx \\ &\geq [\bar{N}_P V_3]^{-1/3} \frac{3}{4} \left[\int_{\Omega - \mathcal{D}(\bar{N}_P)} f^{-1/3}(x) f(x) dx \right] \\ &\quad - o(\bar{N}_P^{-1/3}) \\ &\geq [\bar{N}_P V_3]^{-1/3} \frac{3}{4} \left[\int_{\Omega} f^{2/3}(x) dx - \underline{f}^{-1/3} \bar{f} \int_{\mathcal{D}(\bar{N}_P)} dz \right] \\ &\quad - o(\bar{N}_P^{-1/3}) \\ &= [\bar{N}_P V_3]^{-1/3} \frac{3}{4} \left[\int_{\Omega} f^{2/3}(x) dx \right] - o(\bar{N}_P^{-1/3}). \end{aligned}$$

Multiplying by $\bar{N}_P^{1/3}$ and taking the limit as $\bar{N}_P \rightarrow \infty$, we obtain the claim. ■

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

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

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

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

Proof: Let $\mathbb{E}[D]$ 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

$$\bar{s} + \frac{\mathbb{E}[D]}{v} \leq \frac{m}{\lambda}. \quad (6)$$

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

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

Multiplying both sides by $\bar{N}_P^{1/3}$ and raising to the 3rd power we obtain

$$\bar{N}_P (1 - \varrho)^3 \geq \frac{\lambda^3 [\bar{N}_P^{1/3} \mathbb{E}[Z]]^3}{m^3 v^3}.$$

Applying Little's Law, i.e. $\bar{N}_P = \lambda \bar{W}$, we get

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

Finally, taking the limit as $\varrho \rightarrow 1^-$ we trivially have that $\bar{N}_P \rightarrow \infty$, hence we can apply lemma 4.3 and obtain the claim. ■

V. LIGHT LOAD POLICIES

A. The m -stochastic queue median policy

In this section we briefly describe a policy that achieves asymptotic optimality in the light-load limit. For an instance of the problem, we consider the placement of m depots at the m -stochastic median, i.e., at the locations $x_1^*, \dots, x_m^* = \arg \min_{x_1, \dots, x_m} \mathbb{E}[\min_{i=1}^m \|X - x_i\|]$. Each depot will correspond to a queue, and is assigned a service vehicle.

The m -stochastic queue median policy (SQM)

— 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.

The performance of the SQM policy is characterized by the following theorem.

Theorem 5.1: As $\varrho \rightarrow 0^+$, the SQM policy is asymptotically optimal, that is, $\bar{T}_{\text{SQM}} \rightarrow \bar{T}^*$, as $\varrho \rightarrow 0^+$.

Proof: 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], one can show that the time spent waiting for the vehicle to service other demands goes to zero as $\lambda_i \rightarrow 0^+$, and the system times for the demands serviced by the i th vehicle tends to $\mathbb{E}[\|X - x_i^*\|] / v + \bar{s}_i$, 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 and take $\lambda \rightarrow 0^+$, we find that the expected system time under this policy approaches exactly

$$\bar{T}_{\text{SQM}} \rightarrow \frac{1}{v} \mathbb{E} [\min_{i=1}^m \|X - x_i^*\|] + \bar{s},$$

showing the tightness of the lower bound and the optimality of the SQM policy in the light-load limit. ■

VI. HEAVY LOAD POLICIES

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

A. Bipartite matching tour

Let $X_1, \dots, X_n, Y_1, \dots, Y_n$ be points in \mathbb{R}^d . The points $\mathcal{X} = \{X_1, \dots, X_n\}$ will be pick-up locations and the points $\mathcal{Y} = \{Y_1, \dots, Y_n\}$ will be the corresponding delivery locations. A bipartite matching tour approximates the shortest length tour through the points \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 known in the literature as the *stacker crane tour*).

The bipartite matching tour is constructed as follows. First we add n directed edges $X_i \rightarrow Y_i$ that connect pick-up locations to 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 we obtain one or more (cyclic) tours, which we call *secondary tours*. To connect these, we add the edges of a TSP tour through the locations X_1, \dots, X_n , which we call the *primary tour*. A bipartite matching tour is then as follows: we start at some location in \mathcal{X} (e.g., the closest to the current vehicle's location), say X_j . We service the secondary tour starting at X_j (which returns there). We then continue the primary tour until we find the next unvisited point in \mathcal{X} , say X_k . The procedure is iterated in this way until we reach X_j again (see figure 2). This concept was originally introduced in [18].




Fig. 1. A bipartite matching tour. 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 (starting at P_1) is: $P_1 \rightarrow D_1 \rightarrow P_2 \rightarrow D_2 \rightarrow P_1 \rightarrow P_2 \rightarrow P_3 \rightarrow D_3 \rightarrow P_3 \rightarrow P_1$.

B. The tour splitting policy

In this section we present an *unbiased* service policy, which we call the tour splitting policy (TS).

The tour splitting policy (TS) — Let D be the set of outstanding demands waiting for service. If $D = \emptyset$, the vehicles all move to depot locations. If, instead, $D \neq \emptyset$, compute a bipartite matching tour servicing all demands in D . Split the tour into fragments in the following way: Follow the tour from a randomly chosen starting point. Record the pickup-delivery pairs in the order they are encountered until the distance traveled first exceeds $1/m$ of the tour length (or until the tour is finished). Skip the first demand if it is a delivery site, but always deliver the last demand. Assign the resulting fragment to an available vehicle. To construct additional fragments, repeat the process from where the last fragment ended. Each of the m vehicles services its share of the demands by following its assigned fragment, and then waits. Repeat the process after all vehicles have completed their fragments.

C. Analysis

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

Theorem 6.1 (Performance of TS policy in heavy load): As $\rho \rightarrow 1^-$, the system time for the TS policy satisfies

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

The proof of Theorem 6.1 builds on a number of intermediate results; we start with the following lemma, similar to lemma 1 in [19], characterizing the number of outstanding demands in heavy load.

Lemma 6.2 (Number of demands in heavy load): In heavy load (i.e., $\rho \rightarrow 1^-$), 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 [20] and thus it is omitted. ■

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, $D \neq \emptyset$ with high probability; accordingly, in the following heavy-load analysis we consider the condition $D = \emptyset$ always false.

We refer to the time instant t_i , $i \geq 0$, at which the i th bipartite matching tour is computed, as the epoch i of the policy; we refer to the time interval between epoch i and epoch $i + 1$ as the i th iteration. Let n_i be the number of outstanding demands serviced during iteration i , and let C_i be the length of the i th iteration. Demands arrive according to a Poisson process with rate λ ; then, we have $\mathbb{E}[n_{i+1}] = \lambda \mathbb{E}[C_i]$. The time interval C_i is equal to the maximum of the time intervals needed by the vehicles to service their

assigned shares of demands. Let $T(n_i)$ denote the time to traverse the bipartite matching tour through the n_i demands serviced during iteration i . Then, one can easily upper bound $\mathbb{E}[C_i]$ as

$$\mathbb{E}[C_i] \leq \frac{\mathbb{E}[T(n_i)]}{m} + 2 \frac{D(\Omega)}{v},$$

where $D(\Omega) \doteq \max\{\|p - q\| \mid p, q \in \Omega\}$ is the diameter of Ω . The factor 2 accounts for the travel required to reach the assigned fragment and the possible extra travel required to service the last demand on the assigned fragment. The time to traverse the bipartite matching tour through the n_i demands, i.e., $T(n_i)$, 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 we can express an upper bound for $\mathbb{E}[T(n_i)]$: With a slight abuse of notation, we denote the length of the TSP tour through n_i demands as $L_{\text{TSP}}(n_i)$, and the length of the bipartite matching of n_i delivery to pickup sites as $L_{\text{M}}(n_i)$. We have

$$\begin{aligned} \mathbb{E}[T(n_i)] &= \frac{1}{v} \mathbb{E}[L_{\text{TSP}}(n_i)] + \frac{1}{v} \mathbb{E}[L_{\text{M}}(n_i)] + \mathbb{E}\left[\sum_{j=1}^{n_i} s_j\right] \\ &\leq \mathbb{E}[n_i]^{2/3} \frac{\beta_{\text{TSP}} + \beta_{\text{M}}}{v} \int_{\Omega} f^{2/3}(x) dx + \mathbb{E}[n_i] \bar{s}, \end{aligned} \quad (8)$$

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

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

$$\begin{aligned} \bar{n}_{i+1} &= \lambda \mathbb{E}[C_i] \leq \lambda \left(\frac{\mathbb{E}[T(n_i)]}{m} + 2 \frac{D(\Omega)}{v} \right) \\ &\leq \frac{\lambda}{m} \left(\bar{n}_i^{2/3} \frac{\beta_{\text{TSP}} + \beta_{\text{M}}}{v} \int_{\Omega} f^{2/3}(x) dx + \bar{n}_i \bar{s} + 2m \frac{D(\Omega)}{v} \right). \end{aligned} \quad (9)$$

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

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

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

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 $2mD(\Omega)/v$ is negligible compared to the other terms in the right hand side of equation (9), 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

$$y(i+1) = \frac{\lambda}{m} \left(y(i)^{2/3} \frac{\beta_{\text{TSP}} + \beta_{\text{M}}}{v} \int_{\Omega} f^{2/3}(x) dx + y(i) \bar{s} \right), \quad (10)$$

where $i \geq \bar{i}$ and $y(\bar{i}) = \bar{n}_{\bar{i}}$. System-Y is obtained by replacing the inequality in equation (9) with an equality. Pick, now, any $\varepsilon > 0$. From Young's inequality

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

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 (10) we obtain

$$\begin{aligned} y(i+1) &\leq \frac{\lambda}{m} \left(\bar{s} + \varepsilon \frac{\beta_{\text{TSP}} + \beta_{\text{M}}}{v} \int_{\Omega} f^{2/3}(x) dx \right) y(i) \\ &\quad + \underbrace{\frac{4\lambda/m}{27\varepsilon^2} \frac{\beta_{\text{TSP}} + \beta_{\text{M}}}{v} \int_{\Omega} f^{2/3}(x) dx}_{=O(1)(\lambda/m)\varepsilon^2}. \end{aligned} \quad (11)$$

Next, define System-Z as

$$\begin{aligned} z(i+1) &= \frac{\lambda}{m} \left(\bar{s} + \varepsilon \frac{\beta_{\text{TSP}} + \beta_{\text{M}}}{v} \int_{\Omega} f^{2/3}(x) dx \right) z(i) \\ &\quad + O(1) \frac{\lambda/m}{\varepsilon^2}, \end{aligned} \quad (12)$$

where $i \geq \bar{i}$ and $z(\bar{i}) = \bar{n}_{\bar{i}}$. It is immediate to show that the condition $\bar{n}_i = y(i) = z(i)$ implies that

$$\bar{n}_i \leq y(i) \leq z(i), \quad \text{for all } i \geq \bar{i}. \quad (13)$$

The proof now proceeds as follows. First, we show that the trajectories of System-Z are bounded; this fact, together with equation (13), implies that also trajectories of variables \bar{n}_i and $y(i)$ are bounded. Then, we will compute $\limsup_{i \rightarrow +\infty} y(i)$; this quantity, together with equation (13), will yield the desired result.

Let us consider the first issue, namely boundedness of trajectories for System-Z (described in equation (12)). System-Z is a discrete-time linear system and can be rewritten in compact form as

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

where $\varrho = \lambda \bar{s}/m$ and $b = \lambda(\beta_{\text{TSP}} + \beta_{\text{M}}) \int_{\Omega} f^{2/3}(x) dx/mv$. 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 , the solution $i \mapsto z(i) \in \mathbb{R}_{\geq 0}$ of System-Z converges exponentially fast to the unique equilibrium point

$$z^*(\varepsilon) = (1 - \varrho - \varepsilon b)^{-1} O(1) \frac{\lambda/m}{\varepsilon^2}. \quad (14)$$

Combining equation (13) with the previous statement, we see that the solutions $i \mapsto \bar{n}_i$ and $i \mapsto y(i)$ are bounded. Thus

$$\limsup_{i \rightarrow +\infty} \bar{n}_i \leq \limsup_{i \rightarrow +\infty} y(i) < +\infty. \quad (15)$$

Finally, we turn our attention to the computation of $y \doteq \limsup_{i \rightarrow +\infty} y(i)$. Taking the limsup of the left- and right-hand sides of equation (10), and noting that $\limsup_{i \rightarrow +\infty} y^{2/3}(i) = \left(\limsup_{i \rightarrow +\infty} y(i) \right)^{2/3}$, since $x \rightarrow x^{2/3}$ is continuous and strictly monotone increasing on $\mathbb{R}_{>0}$, we obtain that

$$y = \frac{\lambda}{m} y^{2/3} \frac{\beta_{\text{TSP}} + \beta_{\text{M}}}{v} \int_{\Omega} f^{2/3}(x) dx + y \varrho; \quad (16)$$

rearranging we obtain

$$y = \frac{\lambda^3 (\beta_{\text{TSP}} + \beta_{\text{M}})^3 \left(\int_{\Omega} f^{2/3}(x) dx \right)^3}{m^3 v^3 (1 - \varrho)^3}.$$

Noting that from equation (15) $\limsup_{i \rightarrow +\infty} \bar{n}_i \leq y$, we obtain the desired result. \blacksquare

We are now in a position to prove Theorem 6.1.

Proof: [Proof of Theorem 6.1] Define $\bar{C} \doteq \limsup_{i \rightarrow +\infty} \mathbb{E}[C_i]$; then we have, by using the upper bound on $\mathbb{E}[C_i]$ in equation (9) (and neglecting the constant term $2mD(\Omega)/v$),

$$\begin{aligned} \bar{C} &\doteq \limsup_{i \rightarrow +\infty} \mathbb{E}[C_i] \\ &\leq \frac{1}{m} \left(\bar{n}^{2/3} \frac{\beta_{\text{TSP}} + \beta_{\text{M}}}{v} \int_{\Omega} f^{2/3}(x) dx + \bar{n} \bar{s} \right) \\ &\leq \frac{\lambda^2 (\beta_{\text{TSP}} + \beta_{\text{M}})^3 \left(\int_{\Omega} f^{2/3}(x) dx \right)^3}{v^3 m^3 (1 - \varrho)^2} \\ &\quad + \frac{\lambda^2 \varrho (\beta_{\text{TSP}} + \beta_{\text{M}})^3 \left(\int_{\Omega} f^{2/3}(x) dx \right)^3}{v^3 m^3 (1 - \varrho)^3}. \end{aligned}$$

Hence, in the limit $\varrho \rightarrow 1^-$, we have $\bar{C} \leq \lambda^2 (\beta_{\text{TSP}} + \beta_{\text{M}})^3 \left(\int_{\Omega} f^{2/3}(x) dx \right)^3 / v^3 m^3 (1 - \varrho)^3$.

Since a demand is equally likely to arrive at any moment during its ‘‘arrival’’ iteration, and since is equally likely to be serviced at any sequence in the next iteration (recall that the splitting process of the bipartite matching tour starts at a random point on such a tour), the expected steady-state system time of a random demand, \bar{T} , can be upper bounded, as $\varrho \rightarrow 1^-$, by

$$\bar{T} \leq \frac{1}{2} \bar{C} + \frac{1}{2} \bar{C} \leq \frac{\lambda^2 (\beta_{\text{TSP}} + \beta_{\text{M}})^3 \left(\int_{\Omega} f^{2/3}(x) dx \right)^3}{v^3 m^3 (1 - \varrho)^3}. \quad (17)$$

This completes the proof. \blacksquare

D. Discussion

With theorem 6.1 one can readily prove that the steady-state system time under the TS policy differs from the heavy-load lower bound in theorem 4.4 by a known *constant* factor.

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

$$\frac{\bar{T}_{\text{TS}}}{\bar{T}^*} \leq \frac{(\beta_{\text{TSP}} + \beta_{\text{M}})^3}{\gamma_3^3}, \quad \text{as } \varrho \rightarrow 1^-.$$

Proof: The proof is an immediate consequence of theorems 4.4 and 6.1. \blacksquare

Note that the constant factor in the above theorem is independent of the number of vehicles, the arrival rate of demands, and the spatial density for the demand locations. The system time under the TS policy has the *same* growth rate in terms of ϱ of the lower bound (this is not the case in the work [9]).

The TS policy requires on-line solutions of possibly large TSPs and bipartite matchings. The decision version of the TSP belongs to the class of NP-complete problems, which suggests that there is no general algorithm capable of finding the optimal tour in an amount of time polynomial in the size of the input. Meanwhile, for the (Euclidean) bipartite matching problem there is a known algorithm of order $O(n^{2+\epsilon})$ runtime, where n is the size of the input and ϵ is an arbitrarily small positive constant [21]. In practice, implementations of the TS policy should rely on heuristics or on polynomial-time approximation schemes for the solutions of TSPs, such as Lin-Kernighan’s [22] or Christofides’ [23]; and even bipartite matchings could be computed using the near-linear constant-factor approximation algorithm in [24], whose complexity is $O(n^{1+\epsilon})$.

Finally, we mention that we studied also a different assignment criterion, whereby a newly arrived demand is assigned to vehicle i with probability $1/m$. Interestingly, such ‘‘randomized assignment’’ leads to a heavy-load scaling with respect to the number of vehicles of the order $1/m^2$ (recall that the TS policy has a heavy-load scaling of the order $1/m^3$). Hence, as a general rule, randomized assignment *should be avoided* in heavily congested transportation-on-demand systems.

VII. SIMULATION

In this section we present simulation results for the tour splitting policy. Tours of the TS policy were computed using `linkern`¹ as a solver to generate approximations to the optimal TSP tour. A Python implementation of the Kuhn-Munkres assignment algorithm [25] was used to generate Euclidean bipartite matchings.

In all simulations we assumed the environment Ω to be the unit cube $[0, 1]^3$, with vehicles of unit speed, and the spatial demand density f to be uniform over Ω . For any set of parameters (e.g., ϱ , m etc.) a simulation spanned a window of time $t \in [0, 5000]$. Mean system time was computed, ignoring demands arriving before $t = 3000$ —enough to remove a transient.

In figure 2(a) we show the dependence of the system time \bar{T}_{TS} on the load factor ϱ with a number of vehicles $m = 5$. We consider values of $\varrho \in [0.6, 0.75]$, which correspond to a moderate/heavy load. One can observe that the average system time begins to obey the asymptotic bounds in this

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

regime. We also study how \bar{T}_{TS} scales with m . For a load factor $\rho = 0.65$, we show in figure 2(b) the average system time under TS 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, $\lambda = \rho m / \bar{s}$. Hence, by increasing m while keeping ρ fixed, the lower bound in theorem 4.4 and the upper bound in theorem 6.1 scale as $1/m$. The results show a decreasing trend, which, however, appears to be weaker than the theoretical $1/m$ trend; this is due to one or a combination of the following factors: First, the theoretical results hold formally only in the limit $\rho \rightarrow 1^-$. Second, we are using approximations for optimal TSP and bipartite matching solutions.

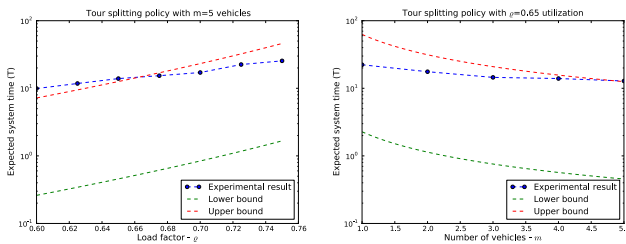


Fig. 2. Performance of TS policy and comparison with upper and lower bounds. Left figure: \bar{T}_{TS} versus ρ . Right figure: scaling of \bar{T}_{TS} with respect to m .

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 TS policy is, asymptotically, within a constant factor of the limit of achievable performance.

This paper leaves numerous important extensions open for further research. Our initial motivation was to study TOD and MOD systems for which f_P and f_D might be drastically different (even with different support), or where a demand's delivery site may not be independent of its pickup. It is of interest also to generalize the results obtained in this paper to vehicles with more complex dynamics, e.g., Dubins vehicles, and to environments with a Manhattan metric. Finally, we plan to consider impatient demands that disappear if they are not serviced within a certain time window.

ACKNOWLEDGMENTS

The authors thank Dr. Ketan Savla for insightful feedback during the development of this work. 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, *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] H. Xu. *Optimal policies for stochastic and dynamic vehicle routing problems*. Dept. of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA, 1995.
- [19] 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.
- [20] M. Pavone, E. Frazzoli, and F. Bullo. Distributed and adaptive algorithms for vehicle routing in a stochastic and dynamic environment. *IEEE Transactions on Automatic Control*, 2009. To Appear, available at <http://arxiv.org/abs/0903.3624>.
- [21] P. K. Agarwal, A. Efrat, and M. Sharir. Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications. In *Proc. of the 11th annual symposium on Computational Geometry*, pages 39–50, Vancouver, Canada, 1995.
- [22] S. Lin and B. W. Kernighan. An effective heuristic algorithm for the traveling-salesman problem. *Operations Research*, 21:498–516, 1973.
- [23] N. Christofides. Worst-case analysis of a new heuristic for the traveling salesman problem. Technical Report 388, Graduate School of Industrial Administration, Carnegie Mellon University, 1976.
- [24] P. Agarwal and K. Varadarajan. A near-linear constant-factor approximation for Euclidean bipartite matching? In *Proc. of the 20th annual symposium on Computational Geometry*, pages 247–252, Brooklyn, USA, 2004.
- [25] J. Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38, March 1957.