

Dynamic Vehicle Routing with Stochastic Time Constraints

Marco Pavone and Emilio Frazzoli

Abstract—In this paper we study a dynamic vehicle routing problem where demands have stochastic deadlines on their waiting times. Specifically, a network of robotic vehicles must service demands whose time of arrival, location and on-site service are stochastic; moreover, once a demand arrives, it remains active for a stochastic amount of time, and then expires. An active demand is successfully serviced when one of the vehicles visits its location before its deadline and provide the required on-site service. The aim is to find the minimum number of vehicles needed to ensure that the steady-state probability that a demand is successfully serviced is larger than a desired value, and to determine the policy the vehicles should execute to ensure that such objective is attained. First, we carefully formulate the problem, and we show its well-posedness by providing some novel ergodic results. Second, we provide a lower bound on the optimal number of vehicles; finally, we analyze two service policies, and we show that one of them is optimal in light load. Simulation results are presented and discussed.

I. INTRODUCTION

The last decade has seen an increasing number of application domains where networks of robotic vehicles are required to visit demands that are spatially *distributed* over an environment, and that possibly require some type of *on-site service*. Surveillance missions by teams of unmanned aerial vehicles are a first clear example; in this case, demands are targets whose potential hazard has to be assessed. Automatic delivery of payloads by mobile robotic networks provides a second example. Other examples include transportation-on-demand systems, automated refuse collection trucks, etc. In many of the above examples, demands arrive *dynamically* in time, and their on-site service is stochastic; moreover, *timeliness* is paramount: should the vehicles take too long to reach the location of the demand, the latter (i) may have escaped and be hard to track (example 1), or (ii) may no longer be interested in the delivered good (example 2). In other words, the routing of robotic vehicles is often dynamic *and* time-constrained, in the sense that demands have (possibly stochastic) deadlines on their waiting times.

Routing problems with both a dynamic and a stochastic component, which are collectively called Dynamic Vehicle Routing Problems (DVRPs), have been extensively studied in the last 20 years [1], [2], [3]; however, little is known about time-constrained versions of DVRPs, despite their practical relevance. The purpose of this paper is to fill this gap. Specifically, we introduce and study the following problem, which we call the Dynamic Vehicle Routing Problem with Stochastic Time Constraints (DVRPSTC): m

vehicles operating in a bounded environment and traveling with bounded velocity must service demands whose time of arrival, location and on-site service are stochastic; moreover, once a demand arrives, it remains active for a stochastic amount of time, and then expires. An active demand is successfully serviced when one of the vehicles visits its location before its deadline and provide the required on-site service. The aim is to find the minimum number of vehicles needed to ensure that the steady-state probability that a demand is successfully serviced is larger than a desired value $\phi^d \in (0, 1)$, and to determine the policy the vehicles should execute to ensure that such objective is attained. Some of the characteristics of the DVRPSTC have been studied in isolation in the literature. When there is *no dynamic* component, and all problem data are known with certainty, the DVRPSTC is closely related to the well-known Vehicle Routing Problem with Time Windows (VRPTW). The VRPTW has been the subject of intensive research efforts for both heuristic and exact optimization approaches (see, e.g., [4] and references therein). When there is *no spatial* component, i.e., all demands arrive at a specific facility, the DVRPSTC becomes a queueing problem with impatient customers, which again has been the subject of intensive study (see, for example, [5], [6]). Finally, the DVRPSTC is also related to coverage problems of mobile sensor networks [7].

The contribution of this paper is threefold: First, we carefully formulate the DVRPSTC. Second, we establish a lower bound on the optimal number of vehicles; in deriving such lower bound, we introduce a novel type of facility location problem, and we provide some analysis and algorithms for it. Third, we analyze two service policies; in particular, we show that one of the two policies is optimal in light load (i.e., when there are few arrivals per unit of time), and we discuss scaling laws for the minimum number of vehicles. The significance of our results stems from two facts: First, the structural insights into the DVRPSTC provide the system designer with essential information to build business and strategic planning models regarding, e.g., fleet sizing and depot locations. Second, our analysis provides directions and guidelines to route the robotic vehicles once the system is deployed. We finally point out that when (i) the arrival rate of demands tends to infinity, (ii) ϕ^d tends to one, and (iii) the deadlines are deterministic, the DVRPSTC reduces to the problem we studied in [8].

II. PRELIMINARIES

A. Notation

We let $\mathbb{R}_{>0}$ and $\mathbb{R}_{\geq 0}$ denote the positive and nonnegative real numbers, respectively. We let $\mathbb{N}_{>0}$ and \mathbb{N} denote the positive and nonnegative integers, respectively. We let $\|\cdot\|$ denote the Euclidean norm in \mathbb{R}^2 . For a measurable and

M. Pavone and E. Frazzoli are with the Laboratory for Information and Decision Systems, Aeronautics and Astronautics Department, Massachusetts Institute of Technology, Cambridge, MA 02139, USA; email: {pavone, frazzoli}@mit.edu.

This research was partially supported by the National Science Foundation, through grants #0705451 and #0705453.

bounded subset $\mathcal{E} \subset \mathbb{R}^2$, we let $|\mathcal{E}|$ denote its area, and we define its diameter as $d_{\mathcal{E}} \doteq \sup\{\|p - q\| \mid p, q \in \mathcal{E}\}$.

B. Asymptotic and Worst-Case Properties of the Traveling Salesman Problem

Given a set D of n points in \mathbb{R}^2 , let $\text{TSP}(D)$ denote the length of the shortest tour through all the points in D ; by convention, $\text{TSP}(\emptyset) = 0$. Assume that the locations of the n points are random variables independently and identically distributed in a compact set \mathcal{E} ; in [9] it is shown that there exists a constant β_{TSP} such that, almost surely, $\lim_{n \rightarrow +\infty} \text{TSP}(D)/\sqrt{n} = \beta_{\text{TSP}} \int_{\mathcal{E}} \sqrt{f(x)} dx$, where f is the density of the absolutely continuous part of the point distribution. The constant β_{TSP} has been estimated numerically as $\beta_{\text{TSP}} \simeq 0.7120 \pm 0.0002$ [10]. Furthermore, for a certain environment \mathcal{E} , the following (deterministic) bound holds on the length of the TSP tour, uniformly on n [9]:

$$\text{TSP}(D) \leq \beta_{\mathcal{E}} \sqrt{|\mathcal{E}|} \sqrt{n}, \quad (1)$$

where $\beta_{\mathcal{E}}$ is a constant that only depends on the shape of \mathcal{E} . We will refer to $\beta_{\mathcal{E}}$ as the characteristic constant of \mathcal{E} .

C. Voronoi Diagrams

Assume that $\mathbf{x} = (x_1, \dots, x_m)$ is an ordered set of m distinct points in $\mathcal{E} \subset \mathbb{R}^2$. The *Voronoi diagram* $\mathcal{V}(\mathbf{x}) = (V_1(\mathbf{x}), \dots, V_m(\mathbf{x}))$ of \mathcal{E} generated by \mathbf{x} is defined by $V_k(\mathbf{x}) = \{x \in \mathcal{E} \mid \|x - x_k\| \leq \|x - x_j\|, \forall j \neq k, j \in \{1, \dots, m\}\}$, $k \in \{1, \dots, m\}$. We refer to \mathbf{x} as the set of *generators* of $\mathcal{V}(\mathbf{x})$, and to $V_k(\mathbf{x})$ as the Voronoi cell of the k th generator. If \mathcal{E} is convex, then each Voronoi cell is a convex set. Lastly, we define $\mathcal{X}_{\text{coinc}} \doteq \{(x_1, \dots, x_m) \in \mathcal{E}^m \mid x_k = x_j \text{ for some } k \neq j\}$.

D. Regenerative Processes

A stochastic process $\{X(t); t \in \mathbb{T}\}$, where $\mathbb{T} = \mathbb{N}$ or $\mathbb{T} = \mathbb{R}_{\geq 0}$, is said to be *regenerative* if it can be split into independent and identically distributed (i.i.d.) *cycles*, or, in other words, if an imbedded renewal process can be found (we refer the reader to see [11, page 169] for a more formal definition). The power of the concept of regenerative processes lies in the existence of a limiting distribution under conditions that are very mild and easy to verify. We next provide some results about discrete-time (i.e. $\mathbb{T} = \mathbb{N}$) regenerative processes. Before proceeding, we need the following definition. A discrete probability distribution f_k , $k \in \mathbb{N}$, is said to be *periodic* if there exists an integer $p > 1$ such that all f_k 's, except, perhaps, $f_p, f_{2p}, f_{3p}, \dots$, vanish. Often, it is easy to check if a distribution is non-periodic.

Lemma 2.1 (Lemma 2 in [12]): A discrete probability distribution f_k , $k \in \mathbb{N}$, is non-periodic if $f_1 > 0$.

Consider a regenerative process $\{X_j; j \in \mathbb{N}\}$, with associated renewal sequence $\{Y_j; j \in \mathbb{N}_{>0}\}$, and let $Y = Y_1$ be the first regeneration cycle; moreover, let $f_k = \mathbb{P}[Y = k]$, $k \in \mathbb{N}$. The stationary version $\{X_j^*; j \in \mathbb{N}\}$ of the regenerative process $\{X_j; j \in \mathbb{N}\}$ is defined by (see [12])

$$\mathbb{P}[X_j^* \in A] = \frac{1}{\mathbb{E}[Y]} \sum_{k=0}^{+\infty} \mathbb{P}[X_{j+k} \in A \mid Y > k] \cdot \mathbb{P}[Y > k], \quad (2)$$

for every $j \in \mathbb{N}$ and every Borel set A .

Theorem 2.2 (Adapted form theorem 2 in [12]): Let $\{X_j; j \in \mathbb{N}\}$ be a nonnegative regenerative process with first regeneration cycle Y , with $\mathbb{E}[Y] < +\infty$. Then, the stationary process given by (2) is well defined and has a proper distribution function, which is independent of j . Moreover,

- (i) $\mathbb{E}[X_0^*] = \mathbb{E}\left[\sum_{j=0}^{Y-1} X_j\right] / \mathbb{E}[Y]$;
- (ii) if f_k , $k \in \mathbb{N}$, is non-periodic, then for every Borel set A , $\lim_{j \rightarrow +\infty} \mathbb{P}[X_j \in A] = \mathbb{P}[X_0^* \in A]$.

III. PROBLEM SETUP

In this section, we set up the problem we wish to study.

A. The Model

Let the workspace $\mathcal{E} \subset \mathbb{R}^2$ be a compact, convex set. A total of m holonomic vehicles operate in \mathcal{E} ; the vehicles are free to move, traveling at a maximum velocity v , within \mathcal{E} . The vehicles are identical, have unlimited range and demand servicing capacity; moreover, each vehicle is associated with a depot whose location is at $x_k \in \mathcal{E}$, $k \in \{1, \dots, m\}$. Demands are generated according to a homogeneous (i.e., time-invariant) spatio-temporal Poisson process, with time intensity $\lambda \in \mathbb{R}_{>0}$, and uniform spatial density over \mathcal{E} . In other words, demands arrive to \mathcal{E} according to a Poisson process with intensity λ , and their locations $\{X_j; j \in \mathbb{N}\}$ are i.i.d. and distributed according to a uniform density whose support is \mathcal{E} ; moreover, the locations are independent of demands' arrival times and of vehicles' positions. Let $\{T_j; j \in \mathbb{N}\}$ denote the sequence of arrival times of demands; we assume that $T_0 = 0$, and that the first arrival finds the system empty. Let $N(t)$, $t \in \mathbb{R}_{\geq 0}$, denote the number of arrivals in $[0, t]$, i.e., $N(t) = \max\{j \in \mathbb{N} \mid T_j \leq t\}$. Each demand j requires a stochastic amount of on-site service time S_j . A vehicle provides on-site service by staying at the demand's location for the entire on-site service time. On-site service is not-interruptible: once a vehicle starts the service, neither the vehicle can interrupt the service nor the demand can leave the system before service completion. We assume that the nonnegative on-site service times $\{S_j; j \in \mathbb{N}\}$ are i.i.d. and generally distributed according to a distribution function $F_S(s)$ with first moment \bar{S} and maximum value $s_{\max} \geq 0$.

Each demand j waits for the beginning of its service no longer than a stochastic *patience time* P_j . We assume that the nonnegative patience times $\{P_j; j \in \mathbb{N}\}$ are i.i.d. and generally distributed according to a distribution function $F_P(p)$ with first moment \bar{P} and maximum value $p_{\max} > 0$; moreover, we assume that $\mathbb{P}[P_j = 0] = 0$, and that the patience times are independent of demands' arrival times, demands' locations, and vehicles' positions. A vehicle can start the on-site service for the j th demand only within the stochastic time window $[T_j, T_j + P_j)$. If the on-site service for the j th demand is not started before the time instant $T_j + P_j$, then the j th demand is considered lost; in other words, such demand leaves the system and never returns. If, instead, the on-site service for the j th demand is started before the time instant $T_j + P_j$, then the demand is considered successfully serviced (recall our assumption that

on-site service is not interruptible); we call such demand a *successful* demand. The waiting time of demand j , denoted by W_j , is the elapsed time between the arrival of demand j and the time either one of the vehicles starts its service or such demand departs from the system due to impatience, whichever happens first. Hence, the j th demand is considered serviced if and only if $W_j < P_j$. Finally, let $\{T_j^s; j \in \mathbb{N}\}$ denote the sequence of arrival times of successful demands; note that the sequence $\{T_j^s; j \in \mathbb{N}\}$ is a thinning of $\{T_j; j \in \mathbb{N}\}$. Let $N^s(t)$, $t \in \mathbb{R}_{\geq 0}$, denote the number of arrivals in $[0, t]$ that will eventually be serviced, i.e., $N^s(t) = \max\{j \in \mathbb{N} | T_j^s \leq t\}$.

An instance of the problem is represented by the data vector $\mathcal{I} = \{\mathcal{E}, v, \lambda, F_S(s), F_P(p)\}$; the number m of vehicles and their routing policy are, instead, decision variables.

B. Information Structure and Control Policies

We first describe the information on which a control policy can rely. We identify four types of information: 1) *Arrival time and location*: we assume that the information on arrivals and locations of demands is immediately available to control policies; 2) *On-site service*: the on-site service requirement of demands may either (i) be available, or (ii) be available only through prior statistics, or (iii) not be available to control policies; 3) *Patience time*: the patience time of demands may either (i) be available, or (ii) be available only through prior statistics; 4) *Departure notification*: the information that a demand leaves the system due to impatience may or may not be available to control policies (if the patience time is available, such information is clearly available).

Hence, we identify a total of nine possible information structures. The *least informative* case is when on-site service requirements and departure notifications are not available, and patience times are available only through prior statistics; the *most informative* case is when on-site service requirements and patience times are available.

We next define the notion of *outstanding demand* for different information structures. If departure notifications are available, a demand is considered outstanding if (i) no vehicle has yet reached its location, (ii) the demand is still in the system. When departure notifications are *not* available, a demand is considered outstanding if (i) no vehicle has yet reached its location, (ii) the elapsed time from its arrival is less than p_{\max} (note that p_{\max} is always known by the vehicles). Note that, in absence of departure notifications, a vehicle will sometimes reach locations of demands that are no longer in the system.

Given an instance \mathcal{I} and a particular information structure (one of the nine possible), let \mathcal{P} be the set of all causal, stationary, and *work-conserving* policies. In this paper, a policy is said to be work-conserving if (i) when a vehicle has no outstanding (in the above sense) demands to service, it moves rectilinearly to (or remains at) its depot location, (ii) when there are outstanding demands, there is at least one vehicle providing service to them (either on-site or by traveling). Property (i) is a technical condition needed to ensure that the underlying stochastic processes are regenerative, while property (ii) is a standard assumption needed to avoid pathological situations. The system is said to be *idling*

if all vehicles are at their depot locations and there are no outstanding demands. We assume that initially all vehicles are at their depots.

We view each policy $\pi \in \mathcal{P}$ as a function of the number m of available vehicles; when needed, we make this dependency explicit by writing $\pi = \pi(m)$. Let $\mathbb{P}_{\pi(m)}[W_j < P_j]$ be the probability, under a policy $\pi(m)$, that the j th demand is serviced. We will show in section IV that under any policy belonging to \mathcal{P} the sequence of acceptance probabilities $\{\mathbb{P}_{\pi(m)}[W_j < P_j]; j \in \mathbb{N}\}$ is convergent. We, then, define the *success factor* of a policy $\pi(m) \in \mathcal{P}$ as $\phi_{\pi(m)} \doteq \lim_{j \rightarrow +\infty} \mathbb{P}_{\pi(m)}[W_j < P_j]$.

C. Problem Definition

Given an instance \mathcal{I} , a particular information structure, and a desired success factor $\phi^d \in (0, 1)$, the problem is to determine a vehicle routing policy π^* that guarantees a success factor at least as large as ϕ^d with the minimum possible number of vehicles. Formally, for a policy $\pi \in \mathcal{P}$, define m_π^* as the solution to the minimization problem

$$\begin{aligned} \min_{m \in \mathbb{N}_{>0}} \quad & m, \\ \text{subject to} \quad & \phi_{\pi(m)} \geq \phi^d. \end{aligned} \quad (3)$$

If the set of feasible solutions is empty, we set, by convention, $m_\pi^* = \infty$. Then, in this paper, we wish to solve the following minimization problem

$$\mathcal{OPT} : \min_{\pi \in \mathcal{P}} m_\pi^*. \quad (4)$$

In principle, one should study the problem \mathcal{OPT} for each of the nine possible information structures. In this paper, instead, we consider the following strategy: first, we derive a lower bound that is valid under the most informative information structure (this implies validity under any information structure), then we present and analyze two service policies that are valid under the least informative information structure (this implies implementability under any information structure). Such approach will give general insights into the problem \mathcal{OPT} .

We start by showing the well-posedness of the problem.

IV. WELL-POSEDNESS

Here, by well-posedness, we mean the existence of a limit for the sequence $\{\mathbb{P}_{\pi(m)}[W_j < P_j]; j \in \mathbb{N}\}$. A demand that finds the system idling faces a situation *probabilistically identical* to that of the first demand. Hence, in our model, all of the relevant stochastic processes are *regenerative*, and the regeneration points are the time instants in which an arrival finds the system idling. With the above discussion in mind, consider the following quantities. Let $\{C_i; i \in \mathbb{N}_{>0}\}$ be the sequence of successive busy cycles: a busy cycle is the duration between two successive arrival epochs of demands finding the system idling. The C_i 's are i.i.d. random variables on $\mathbb{R}_{\geq 0}$. Let $\{B_i; i \in \mathbb{N}_{>0}\}$ be the sequence of successive busy periods: the busy period is the part of the busy cycle during which at least one vehicle is providing service (either by traveling or on-site) to a demand, or it is moving to its depot. The B_i 's are i.i.d. random variables on $\mathbb{R}_{\geq 0}$. Let $\{L_i; i \in \mathbb{N}_{>0}\}$ (or $\{L_i^s; i \in \mathbb{N}_{>0}\}$) be the number

of demands arrived in the system (or successfully serviced) during the i th busy period, including the one initializing it. The L_i 's (or L_i^s 's) are i.i.d. random variables on \mathbb{N} . In what follows, we use the imbedded renewal process:

$$\begin{cases} \Lambda_0 = 0, \\ \Lambda_i = \Lambda_{i-1} + L_i, \quad i \geq 1. \end{cases} \quad (\text{on } \mathbb{N})$$

In order to apply the results on regenerative processes, we first have to prove the finiteness of busy cycles.

Lemma 4.1 (Finiteness of busy cycles): Given an instance \mathcal{I} , an information structure, and a policy belonging to \mathcal{P} , we have $\mathbb{E}[C_1] < +\infty$.

Proof: From the definitions we have $\mathbb{E}[C_1] = \mathbb{E}[B_1] + (1/\lambda)$, where B_1 is the part of the first regeneration cycle during which the vehicles continuously work, i.e., there is at least one outstanding demand, or there is at least one vehicle moving to its depot. Note that, if there are no arrivals during a time interval of length $p_{\max} + s_{\max} + d_{\mathcal{E}}/v$, all vehicles will surely be at their depots after that time interval. Hence, we obtain the inequality, for $b \in \mathbb{R}_{\geq 0}$: $\mathbb{P}[B_1 \geq b + p_{\max} + s_{\max} + \frac{d_{\mathcal{E}}}{v}] \leq \mathbb{P}[\text{at least 1 arrival in } [b, b + p_{\max} + s_{\max} + \frac{d_{\mathcal{E}}}{v}]] \mathbb{P}[B_1 \geq b] = (1 - \exp(-\lambda(p_{\max} + s_{\max} + \frac{d_{\mathcal{E}}}{v}))) \mathbb{P}[B_1 \geq b]$. Then, it is easy to show that $\mathbb{E}[B_1] = \int_0^{+\infty} \mathbb{P}[B_1 \geq b] db < +\infty$. Thus, we have $\mathbb{E}[C_1] < +\infty$. ■

A simple relation between C_1 and L_1 is provided by the following lemma.

Lemma 4.2: Given an instance \mathcal{I} , a particular information structure, and a policy belonging to \mathcal{P} , we have $\mathbb{E}[L_1] = \lambda \mathbb{E}[C_1]$.

Proof: The proof is a consequence of Wald's lemma, and it is omitted in the interest of brevity (see section 3.3 in [5] for a similar result). ■

Since the busy cycles are finite, we can use the theory of regenerative processes to prove the well-posedness of the problem.

Theorem 4.3 (Well-posedness): Given an instance \mathcal{I} , a particular information structure, and a policy π belonging to \mathcal{P} , the sequence $\{\mathbb{P}_{\pi(m)}[W_j < P_j]; j \in \mathbb{N}\}$ is convergent, and its limit is equal to $\mathbb{E}[L_1^s]/\mathbb{E}[L_1]$.

Proof: In this proof, to keep the notation simple, we avoid the usage of the subscript $\pi(m)$. Let I_j^s be the indicator random variable

$$I_j^s = \begin{cases} 1 & \text{if } W_j < P_j, \\ 0 & \text{if } W_j = P_j, \end{cases}$$

i.e., I_j^s equals one if the j th demand is successfully serviced. From the previous discussion, the stochastic process $\{I_j^s; j \in \mathbb{N}\}$ is regenerative relative to the discrete-time renewal process $\{\Lambda_i; i \in \mathbb{N}\}$, and it is nonnegative. By lemma 4.1 and lemma 4.2, the expectation of L_1 , which is the first regeneration cycle, is finite. Moreover, it clearly holds $\mathbb{P}[L_1 = 1] > 0$, hence by lemma 2.1 the distribution of L_1 is non-periodic. Let $\{I_j^{s,*}; j \in \mathbb{N}\}$ be the stationary version of $\{I_j^s; j \in \mathbb{N}\}$. Then, by applying theorem 2.2 part (ii), by noting that $\tilde{I}_0^{s,*}$ is an indicator random variable and thus $\mathbb{P}[\tilde{I}_0^{s,*} = 1] = \mathbb{E}[I_0^{s,*}]$, and by finally applying theorem 2.2 part (i), we obtain the series of equalities $\lim_{j \rightarrow +\infty} \mathbb{P}[I_j^s = 1] = \mathbb{P}[\tilde{I}_0^{s,*} = 1] = \mathbb{E}[I_0^{s,*}] =$

$\mathbb{E}\left[\sum_{j=0}^{L_1-1} I_j^s\right]/\mathbb{E}[L_1]$; since $\mathbb{E}\left[\sum_{j=0}^{L_1-1} I_j^s\right] = \mathbb{E}[L_1^s]$, and $\lim_{j \rightarrow +\infty} \mathbb{P}[I_j^s = 1] = \lim_{j \rightarrow +\infty} \mathbb{P}[W_j < P_j]$, we get the claim. ■

It is natural to wonder if the optimization problem \mathcal{OPT} can be restated in terms of time averages, in other words, if the equality

$$\lim_{j \rightarrow \infty} \mathbb{P}_{\pi(m)}[W_j < P_j] \stackrel{?}{=} \lim_{t \rightarrow +\infty} \frac{N^s(t)}{N(t)}$$

holds almost surely. The answer is affirmative, and its proof relies on some arguments in the theory of continuous-time regenerative processes (we omit the details for brevity). The usefulness of such ergodic result stems from two facts: (i) on a theoretical level, formulating the problem in terms of time averages or limiting probabilities is *equivalent*, (ii) on a practical level, in some cases it might be easier to characterize $\lim_{t \rightarrow +\infty} \frac{N^s(t)}{N(t)}$, while in other cases it might be easier to study $\lim_{j \rightarrow \infty} \mathbb{P}_{\pi(m)}[W_j < P_j]$.

V. LOWER BOUND

In this section, we present a lower bound for the optimization problem \mathcal{OPT} that holds under any information structure. This lower bound is intimately related to a novel type of facility location problem, for which we will provide some analysis and algorithms later in this section.

A. Lower Bound

Let $\mathbf{x} = (x_1, \dots, x_m)$ and define

$$H_m(\mathbf{x}) \doteq \frac{1}{|\mathcal{E}|} \int_{\mathcal{E}} \left(1 - F_P \left(\min_{k \in \{1, \dots, m\}} \frac{\|x - x_k\|}{v} \right) \right) dx.$$

Theorem 5.1: Given an instance \mathcal{I} , an information structure, and a desired success factor $\phi^d \in (0, 1)$, the solution to the minimization problem \mathcal{OPT} is lower bounded by the solution to the minimization problem

$$\begin{aligned} & \min_{m \in \mathbb{N}_{>0}} m \\ & \text{subject to} \quad \sup_{\mathbf{x} \in \mathcal{E}^m} H_m(\mathbf{x}) \geq \phi^d. \end{aligned} \quad (5)$$

Proof: Consider a policy $\pi(m) \in \mathcal{P}$, and assume that m vehicles execute such policy. In the remainder of the proof, to keep the notation simple, we avoid the usage of the subscript $\pi(m)$. Consider the j th demand, and let \bar{X}_k be the position of the k th vehicle when such demand arrives. Let $\bar{\mathbf{X}} = (\bar{X}_1, \dots, \bar{X}_m)$. Obviously, the waiting time of demand j is at least as large as the minimum travel time between its position and the closest vehicle's position, i.e., $W_j \geq \min_{k \in \{1, \dots, m\}} \|X_j - \bar{X}_k\|/v$. The vehicles are located in the workspace according to some generally unknown cumulative distribution function that depends on the policy; we denote such distribution function as $F : \mathcal{E}^m \rightarrow [0, 1]$. Then, the acceptance probability for demand j can be bounded according to (recall that X_j and P_j are, by assumption, independent of $\bar{\mathbf{X}}$)

$$\begin{aligned} \mathbb{P}[W_j < P_j] & \leq \mathbb{P}[\min_{k \in \{1, \dots, m\}} \frac{\|X_j - \bar{X}_k\|}{v} < P_j] \\ & = \int_{\mathcal{E}^m} \mathbb{P}[\min_{k \in \{1, \dots, m\}} \frac{\|X_j - \bar{X}_k\|}{v} < P_j | \bar{\mathbf{X}} = \bar{\mathbf{x}}] dF(\bar{\mathbf{x}}) \\ & \leq \int_{\mathcal{E}^m} \sup_{\bar{\mathbf{x}} \in \mathcal{E}^m} \mathbb{P}[\min_{k \in \{1, \dots, m\}} \frac{\|X_j - \bar{x}_k\|}{v} < P_j] dF(\bar{\mathbf{x}}) \end{aligned}$$

$$\begin{aligned}
&= \sup_{\bar{\mathbf{x}} \in \mathcal{E}^m} \mathbb{P}[\min_{k \in \{1, \dots, m\}} \frac{\|X_j - \bar{x}_k\|}{v} < P_j] \\
&= \sup_{\bar{\mathbf{x}} \in \mathcal{E}^m} \frac{1}{|\mathcal{E}|} \int_{\mathcal{E}} \mathbb{P}[\min_{k \in \{1, \dots, m\}} \frac{\|X_j - \bar{x}_k\|}{v} < P_j | X_j = x_j] dx_j \\
&= \sup_{\bar{\mathbf{x}} \in \mathcal{E}^m} \frac{1}{|\mathcal{E}|} \int_{\mathcal{E}} \left(1 - F_P \left(\min_{k \in \{1, \dots, m\}} \frac{\|x_j - \bar{x}_k\|}{v} \right) \right) dx_j.
\end{aligned}$$

Hence, we have $\phi = \lim_{j \rightarrow +\infty} \mathbb{P}[W_j < P_j] \leq \sup_{\bar{\mathbf{x}} \in \mathcal{E}^m} H_m(\mathbf{x})$, and a necessary condition for ϕ to be at least as large as ϕ^d is that $\sup_{\bar{\mathbf{x}} \in \mathcal{E}^m} H_m(\mathbf{x}) \geq \phi^d$. ■

In section VI, we will present a policy that requires the least amount of information and is optimal in light load, i.e., in the limit $\lambda \rightarrow 0^+$.

As a matter of fact, in equation (5) we have implicitly introduced a novel type of facility location problem, which is worth a definition.

Definition 5.2 (The m -LPIC): Given a compact, convex set $\mathcal{E} \subset \mathbb{R}^2$, a cumulative distribution function $F_P : \mathbb{R} \rightarrow [0, 1]$, a constant $v > 0$, and an integer $m \in \mathbb{N}_{>0}$, the m -Location Problem with Impatient Customers (m -LPIC) is the optimization problem: $H^*(m) \doteq \sup_{\mathbf{x} \in \mathcal{E}^m} H_m(\mathbf{x})$.

B. Analysis and Algorithms for the m -LPIC

In this section we study in some detail the m -LPIC. In particular, we study (i) conditions under which a maximizer exists, and (ii) a gradient-ascent law for the optimization of H_m . We begin with the following theorem, which shows that a maximizer for the m -LPIC exists in most practical scenarios.

Theorem 5.3 (Existence of a maximizer): Assume that F_P is piecewise differentiable on $\mathbb{R}_{\geq 0}$ with a finite number of (jump) discontinuities; then, $H_m(\mathbf{x})$ has a global maximum.

Proof: As a consequence of theorem 2.2 in [13], $H_m(\mathbf{x})$ is globally Lipschitz on \mathcal{E}^m . Therefore, $H_m(\mathbf{x})$ is continuous on a compact set (since \mathcal{E} is compact), and, by the extreme value theorem, it has a global maximum. ■

It is also possible to state a differentiability result, which will be the basis of a gradient-ascent algorithm for the optimization of H_m .

Theorem 5.4 (Differentiability of $H_m(\mathbf{x})$): Assume that F_P is differentiable on $\mathbb{R}_{\geq 0}$ with derivative equal to f_P ; then $H_m(\mathbf{x})$ is continuously differentiable on $\mathcal{E}^m \setminus \mathcal{X}_{\text{coinc}}$, where for each $j \in \{1, \dots, m\}$

$$\frac{\partial H_m}{\partial x_j}(\mathbf{x}) = \frac{1}{v|\mathcal{E}|} \int_{V_j(\mathbf{x})} f_P(\|x - x_j\|/v) \frac{x - x_j}{\|x - x_j\|} dx,$$

where $\mathcal{V}(\mathbf{x}) = (V_1(\mathbf{x}), \dots, V_m(\mathbf{x}))$ is the Voronoi diagram generated by $\mathbf{x} = (x_1, \dots, x_m)$.

Proof: One can easily show that part (ii) of theorem 2.2 in [13] is applicable, and the claim is an immediate consequence. ■

Remark 5.5: By using the results in part (ii) of theorem 2.2 in [13], theorem 5.6 can be extended to the case where F_P is piecewise differentiable on $\mathbb{R}_{\geq 0}$ with a finite number of (jump) discontinuities; however, the expression for the gradient is quite cumbersome and is omitted.

By using theorem 5.4 we can readily set up a gradient-ascent law to maximize the locational optimization function

H_m . Specifically, assume that F_P is differentiable on $\mathbb{R}_{\geq 0}$; then, consider the following continuous gradient-ascent law defined over the set \mathcal{E}^m (a discrete version can be similarly stated):

$$\dot{x}_j(t) = \begin{cases} \frac{\partial H_m}{\partial x_j}(\mathbf{x}(t)) & \text{if } \mathbf{x} \in \mathcal{E}^m \setminus \mathcal{X}_{\text{coinc}} \text{ and } x_j \in \text{int}(\mathcal{E}), \\ \text{pr}_{\mathcal{E}} \left(\frac{\partial H_m}{\partial x_j}(\mathbf{x}(t)) \right) & \text{if } \mathbf{x} \in \mathcal{E}^m \setminus \mathcal{X}_{\text{coinc}} \text{ and } x_j \in \partial\mathcal{E}, \\ 0 & \text{otherwise;} \end{cases} \quad (6)$$

where $t \in \mathbb{R}_{\geq 0}$, $j \in \{1, \dots, m\}$, the dot represents differentiation with respect to t , $\text{int}(\mathcal{E})$ is the interior of \mathcal{E} , $\text{pr}_{\mathcal{E}}(\partial H_m / \partial x_j)$ is the orthogonal projection onto \mathcal{E} of $\partial H_m / \partial x_j$, and we assume that the Voronoi diagram is continuously updated. The vector field is discontinuous, so we understand the solutions in the Krasovskii sense; see [14]. The convergence properties of the gradient-ascent law in equation (6) are summarized by the following theorem.

Theorem 5.6 (Convergence of gradient ascent (6)):

Assume that F_P is differentiable on $\mathbb{R}_{\geq 0}$; then, for each initial condition $\mathbf{x}(0) \in \mathcal{E}^m \setminus \mathcal{X}_{\text{coinc}}$, the Krasovskii solution that exactly satisfies (6) monotonically optimizes $H_m(\mathbf{x})$ and asymptotically converges to the union of $\mathcal{X}_{\text{coinc}}$ and the set of critical points of $H_m(\mathbf{x})$.

Proof: The proof of this theorem is very similar to the proof of theorem 4.1 in [15], and it is omitted in the interest of brevity. ■

A possible variant of the gradient ascent (6) consists in setting $\dot{x}_j = 0$, when $\mathbf{x} \in \mathcal{X}_{\text{coinc}}$, only for the points that are co-located; if the points are co-located, the Voronoi diagram is computed by considering the co-located points as a single point. Such variant is still guaranteed to asymptotically converge to the union of $\mathcal{X}_{\text{coinc}}$ and the set of critical points of $H_m(\mathbf{x})$; moreover, it is amenable to distributed implementation, since the gradient ascent law is then *distributed* over the dual of the Voronoi diagram, i.e., over the Delaunay graph. This last feature is especially useful when a large network of robotic vehicles is employed.

VI. AN OPTIMAL LIGHT LOAD POLICY

In this section we propose and analyze a policy that requires the least amount of information and is optimal in light load; this result holds for any instance \mathcal{I} in which F_P satisfies the (mild) assumptions of theorem 5.3.

A. The Policy

The Nearest-Depot Assignment (NDA) policy is described next (note that this policy only requires the knowledge of F_P ; moreover, it is required that $H_m(\mathbf{x})$ has a global maximum).

Nearest-Depot Assignment (NDA) Policy: Let $\bar{\mathbf{x}} \doteq \arg \max_{\mathbf{x} \in \mathcal{E}^m} H_m(\mathbf{x})$ (if there are multiple maxima, pick one arbitrarily), and let \bar{x}_k be the location of the depot for the k th vehicle, $k \in \{1, \dots, m\}$. Assign a newly arrived demand to the vehicle whose depot is the nearest to that demand's location, and let \mathcal{S}_k be the set of outstanding (in the sense of section III-B) demands assigned to vehicle k . If the set \mathcal{S}_k is empty, move to \bar{x}_k ; otherwise, visit demands in \mathcal{S}_k in first-come, first-served order, by taking the shortest path to each demand location. Repeat.

Next theorem shows on a theoretical level the optimality of the NDA policy in light load, and on a practical level how to compute m_{NDA}^* for low values of λ .

Theorem 6.1 (Optimality of NDA policy): Consider an instance \mathcal{I} with F_P satisfying the assumptions of theorem 5.3, and any of the nine possible information structures. Then, for almost all values of $\phi^d \in (0, 1)$ (i.e., except for a set of measure zero) the NDA policy is optimal in light load, i.e.,

$$\limsup_{\lambda \rightarrow 0^+} \frac{m_{\text{NDA}}^*(\lambda)}{\min_{\pi \in \mathcal{P}(\lambda)} m_{\pi}^*(\lambda)} = 1.$$

Proof: Define the countable set $\mathcal{H} \doteq \{H^*(m) \mid m \in \mathbb{N}_{>0}\}$; by definition of H^* , we have $\inf \mathcal{H} \geq 0$, and $\sup \mathcal{H} \leq 1$. Consider any desired success factor $\phi^d \in (0, 1) \setminus \mathcal{H}$; note that \mathcal{H} is a countable set, hence its measure is zero (in other words, we are leaving out a zero-measure set of possible success factors).

Assume that \underline{m} vehicles execute the NDA policy, where \underline{m} is the solution to the minimization problem (5). (Note that \underline{m} is independent of λ .) Consider any $\lambda \in \mathbb{R}_{>0}$, and define the event: $A_j \doteq \{\text{all vehicles at their depots at the arrival epoch of the } j\text{th demand}\}$; the acceptance probability for demand j satisfies the inequality: $\mathbb{P}[W_j < P_j] \geq \mathbb{P}[W_j < P_j \mid A_j] \cdot \mathbb{P}[A_j]$. Assume that $T_j - T_{j-1}$, $j \in \mathbb{N}_{>0}$, is larger than $p_{\max} + s_{\max} + d_{\mathcal{E}}/v$; then, the j th demand finds, surely, all vehicles idling at their depots. Hence, we can lower bound $\mathbb{P}[A_j]$, $j \in \mathbb{N}_{>0}$, according to: $\mathbb{P}[A_j] \geq \mathbb{P}[T_j - T_{j-1} > p_{\max} + s_{\max} + d_{\mathcal{E}}/v] = \exp(-\lambda(p_{\max} + s_{\max} + d_{\mathcal{E}}/v))$; note that this bound is independent of j . Conditioning on the event A_j , all vehicles are at their depots, and therefore $\mathbb{P}[W_j < P_j \mid A_j] = \mathbb{P}[\min_{k \in \{1, \dots, \underline{m}\}} \frac{\|X_j - \bar{x}_k\|}{v} < P_j] = H_{\underline{m}}(\bar{\mathbf{x}}) = H^*(\underline{m})$. Hence, we obtain, for every $\lambda \in \mathbb{R}_{>0}$, $\phi_{\text{NDA}}(\underline{m}) = \lim_{j \rightarrow +\infty} \mathbb{P}[W_j < P_j] \geq H^*(\underline{m}) \exp(-\lambda(p_{\max} + s_{\max} + d_{\mathcal{E}}/v))$. From the definition of \underline{m} (see equation (5)), and from the fact that $\phi^d \in (0, 1) \setminus \mathcal{H}$ (hence $H^*(\underline{m}) = \phi^d$ is impossible), it follows that $H^*(\underline{m}) > \phi^d$. Thus, we conclude that there exists $\Lambda > 0$ such that $\phi_{\text{NDA}}(\underline{m}) \geq \phi^d$ for all $\lambda < \Lambda$. Therefore, there exists $\Lambda > 0$ such that $\underline{m} \geq m_{\text{NDA}}^*(\lambda)$ for all $\lambda < \Lambda$; hence, by applying theorem 5.1, we obtain $\limsup_{\lambda \rightarrow 0^+} \frac{m_{\text{NDA}}^*(\lambda)}{\min_{\pi \in \mathcal{P}(\lambda)} m_{\pi}^*(\lambda)} \leq \limsup_{\lambda \rightarrow 0^+} \frac{\underline{m}}{\underline{m}} = 1$. This concludes the proof. ■

B. Discussion and Simulations

It is natural to wonder if $\phi^d \rightarrow 1^-$ implies $\min_{\pi \in \mathcal{P}} m_{\pi}^* \rightarrow +\infty$; we now have the tools to show that, in general, this is not the case. Consider any of the nine possible information structures; moreover, let \mathcal{I} be an instance where F_P satisfies the assumptions of theorem 5.3, and the support of F_P is $[d_{\mathcal{E}}/v, +\infty)$. It is easy to see that $H^*(1) = 1$; then, by using the same arguments as those in the proof of theorem 6.1, one can show that there exists $\Lambda > 0$ such that for all $\lambda < \Lambda$ it holds $\min_{\pi \in \mathcal{P}} m_{\pi}^* = 1$, for any $\phi^d \in (0, 1)$. This example shows that, in general, $\phi^d \rightarrow 1^-$ does not imply that $\min_{\pi \in \mathcal{P}} m_{\pi}^* \rightarrow +\infty$.

We next provide some simulation results for the NDA policy. We consider patience times that are uniformly distributed in the interval $[0, 1.6]$; moreover, the arrival rate




Fig. 1. Left Figure: Approximate values for $H^*(m)$ (the bars indicate the range of values obtained by maximizing H_m). Right Figure: Experimental values of ϕ_{NDA} . The desired success factor is $\phi^d = 0.9$.

is $\lambda = 5$, the workspace is the unit square, the vehicles' velocity is $v = 1$, and $s_{\max} = 0$ (i.e., there is no on-site service requirement). Finally, we consider a desired success factor $\phi^d = 0.9$. To find a lower bound on the required number of vehicles, we solve the optimization problem (5); in particular, starting from $m = 1$, we compute $H^*(m)$ until $H^*(m) \geq \phi^d$. The solution to the m -LPIC, i.e., the value $H^*(m)$, is approximately computed for each m by performing the gradient-ascent law (6) starting from 10 random initial conditions. In figure 1, the left figure shows the range of values that are obtained by maximizing H_m , for several values of m . It can be noted that for each m the range of values is rather small, in other words the function H_m appears to have maxima whose values are close to each other. From the left figure we estimate (recall that we are using approximate values for $H^*(m)$) a lower bound on the required number of vehicles equal to 7. The right figure shows experimental values of ϕ_{NDA} as a function of the number of agents m . It can be noted that the minimum number of vehicles required by the NDA policy to ensure a success factor at least as large as ϕ^d is 8, in almost perfect accordance with theorem 6.1 (recall that theorem 6.1 formally holds only in the limit $\lambda \rightarrow 0^+$).

VII. A POLICY FOR MODERATE AND HEAVY LOADS

In this section we propose and analyze a policy that is well-defined for any information structure and for any instance \mathcal{I} , however it is particularly tailored for the *least* informative case and is most effective in moderate and heavy loads. The Batch (B) policy is described next.

Batch (B) Policy: Partition \mathcal{E} into m equal area regions \mathcal{E}_k , $k \in \{1, \dots, m\}$, and assign one vehicle to each region. Assign a newly arrived demand that falls in \mathcal{E}_k to the vehicle responsible for region k , and let \mathcal{S}_k be the set of locations of outstanding (in the sense of section III-B) demands assigned to vehicle k . For each vehicle-region pair k : if the set \mathcal{S}_k is empty, move to the median (the ‘‘depot’’) of \mathcal{E}_k ; otherwise, compute a TSP tour through all demands in \mathcal{S}_k and vehicle's current position, and service demands by following the TSP tour, skipping demands that are no longer outstanding. Repeat.

A. Analysis of the Policy

The following theorem characterizes the batch policy, under the assumption $s_{\max} = 0$, and assuming the least informative information structure.

Theorem 7.1 (Vehicles required by batch policy): Given an instance \mathcal{I} with $s_{\max} = 0$, the least informative information structure, and a desired success factor

$\phi^d \in (0, 1)$, the solution for the batch policy to the minimization problem in equation (3) is upper bounded by

$$\tilde{m} \doteq \min \left\{ m \mid \sup_{\theta \in \mathbb{R}_{>0}} (1 - F_P(\theta))(1 - 2g(m)/\theta) \geq \phi^d \right\}, \quad (7)$$

where $g(m) \doteq \frac{1}{2} \left(\frac{\bar{\beta}^2}{v^2} |\mathcal{E}| \frac{\lambda}{m^2} + \sqrt{\frac{\bar{\beta}^4}{v^4} |\mathcal{E}|^2 \frac{\lambda^2}{m^4} + 8 \frac{\bar{\beta}^2}{v^2} |\mathcal{E}| \frac{1}{m}} \right)$, and where $\bar{\beta}$ is a constant that depends on the *shape* of the service regions. In other words, the batch policy with a number of vehicles that satisfies equation (7) guarantees a success factor at least as large as the desired one, i.e., ϕ^d .

Proof: In the batch policy each region has equal area, and contains a single vehicle. Thus, the m vehicle problem in a workspace of area $|\mathcal{E}|$ has been turned into m independent single-vehicle problems, each in a region of area $|\mathcal{E}|/m$, and with Poisson arrivals with rates λ/m . In particular, the well-posedness theorem 4.3 holds within each region. The strategy of the proof is then as follows: assuming that m vehicles execute the batch policy, in part 1) we lower bound the limiting acceptance probability within each region; in other words, we lower bound $\lim_{j_k \rightarrow +\infty} \mathbb{P}[W_{j_k} < P_{j_k}]$, where j_k is the j th demand that falls in region k . Then, in part 2), we lower bound the limiting acceptance probability within the entire workspace, and we conclude the proof.

Part 1): Acceptance probability within a region. Consider a region k , $k \in \{1, \dots, m\}$. For simplicity of notation, we shall omit the label k (e.g., instead of j_k , we simply use j to denote the j th demand that falls in region k). We refer to the time instant in which the vehicle assigned to the region computes the r th, $r \in \mathbb{N}_{\geq 0}$, TSP tour as the epoch r of the policy; for $r \in \mathbb{N}_{>0}$, we refer to the time interval between epoch $(r-1)$ and the time instant in which the vehicle visits the last demand along the $(r-1)$ th TSP tour (possibly skipping some demands) as the r th busy period, and we denote its length with B_r ; similarly, we refer to the time interval between epoch $(r-1)$ and epoch r as the r th busy cycle. Let n_r , $r \in \mathbb{N}_{>0}$, be the number of demands arrived during the r th busy period; we let $n_0 = 0$. The number of demands' locations visited during the $(r+1)$ th busy period, $r \in \mathbb{N}_{\geq 0}$, is almost surely no larger than $\max(n_r, 1)$; in particular, it may happen that during the r th busy period there are no arrivals, and thus the vehicle waits for a new demand and immediately provides service to it (recall, also, that the arrival process to each region is Poisson, and thus the probability of "bulk" arrivals is zero). Define $\bar{\beta} \doteq \max_{k \in \{1, \dots, m\}} \beta_{\mathcal{E}, k}$, where $\beta_{\mathcal{E}, k}$ is the characteristic constant of region k ; by the deterministic inequality for a TSP tour through n points (see equation (1)), we have (recall that the area of the region is $|\mathcal{E}|/m$, and that $s_{\max} = 0$)

$$B_{r+1} \leq \frac{\bar{\beta}}{v} \sqrt{\frac{|\mathcal{E}|}{m}} \sqrt{\max(n_r, 1) + 1}, \quad \text{almost surely}; \quad (8)$$

the $+1$ is needed to take into consideration the vehicle's starting position. By simple inductive arguments, it is immediate to show that both $\mathbb{E}[n_r]$ and $\mathbb{E}[\sqrt{n_r}]$ are finite; hence, by taking expectation in equation (8), and by applying Jensen's inequality for concave functions in the form $\mathbb{E}[\sqrt{X}] \leq \sqrt{\mathbb{E}[X]}$, we get $\mathbb{E}[B_{r+1}] \leq \frac{\bar{\beta}}{v} \sqrt{\frac{|\mathcal{E}|}{m}} \sqrt{\mathbb{E}[\max(n_r, 1)] + 1} \leq \frac{\bar{\beta}}{v} \sqrt{\frac{|\mathcal{E}|}{m}} \sqrt{\mathbb{E}[n_r] + 2}$. By

applying the law of iterated expectation, it is easy to show that the expected number of demands that arrive in the region during the r th busy period, i.e., $\mathbb{E}[n_r]$, is equal to $(\lambda/m) \mathbb{E}[B_r]$. Then, we obtain the following recurrence relation $\mathbb{E}[B_{r+1}] \leq \frac{\bar{\beta}}{v} \sqrt{\frac{|\mathcal{E}|}{m}} \sqrt{\frac{\lambda}{m}} \mathbb{E}[B_r] + 2$. This recurrence relation allows to find an upper bound on $\limsup_{r \rightarrow +\infty} \mathbb{E}[B_r]$; indeed, it is easy to show that $\limsup_{r \rightarrow +\infty} \mathbb{E}[B_r] \leq g(m)$. We are now in a position to lower bound the limiting acceptance probability in region k . Consider, in steady state, a random tagged demand; let \hat{W} be its waiting time, and \hat{P} be its patience time. Moreover, let \hat{R} be the epoch that immediately follows the arrival of the tagged demand. By the law of total probability, we have, for any $\theta \in \mathbb{R}_{>0}$, $\mathbb{P}[\hat{W} < \hat{P}] \geq \mathbb{P}[\hat{W} < \hat{P} | B_{\hat{R}} + B_{\hat{R}+1} < \theta] \mathbb{P}[B_{\hat{R}} + B_{\hat{R}+1} < \theta]$. Since, from the definition of the batch policy, $\hat{W} \leq B_{\hat{R}} + B_{\hat{R}+1}$ surely, we have $\mathbb{P}[\hat{W} < \hat{P} | B_{\hat{R}} + B_{\hat{R}+1} < \theta] \geq \mathbb{P}[\theta < \hat{P} | B_{\hat{R}} + B_{\hat{R}+1} < \theta] = \mathbb{P}[\theta < \hat{P}] = 1 - F_P(\theta)$; in the previous chain of inequalities, the removal of the conditioning on the event $\{B_{\hat{R}} + B_{\hat{R}+1} < \theta\}$ is possible since, under the least informative information structure, the value of $B_{\hat{R}} + B_{\hat{R}+1}$ does not provide *any* information on the value of \hat{P} . Then, by collecting the previous results and applying Markov's inequality, we obtain $\mathbb{P}[\hat{W} < \hat{P}] \geq (1 - F_P(\theta)) \mathbb{P}[B_{\hat{R}} + B_{\hat{R}+1} < \theta] \geq (1 - F_P(\theta))(1 - (\mathbb{E}[B_{\hat{R}}] + \mathbb{E}[B_{\hat{R}+1}])/\theta) \geq (1 - F_P(\theta))(1 - 2g(m)/\theta)$. Since the previous chain of inequalities holds for all $\theta \in \mathbb{R}_{>0}$, we obtain $\mathbb{P}[\hat{W} < \hat{P}] \geq \sup_{\theta \in \mathbb{R}_{>0}} (1 - F_P(\theta))(1 - 2g(m)/\theta)$. Hence, we conclude that within region k it holds $\lim_{j \rightarrow +\infty} \mathbb{P}[W_j < P_j] \geq \sup_{\theta \in \mathbb{R}_{>0}} (1 - F_P(\theta))(1 - 2g(m)/\theta)$.

Part 2): Acceptance probability within the entire workspace. From part 1), we have $\lim_{j_k \rightarrow +\infty} \mathbb{P}[W_{j_k} < P_{j_k}] \geq \sup_{\theta \in \mathbb{R}_{>0}} (1 - F_P(\theta))(1 - 2g(m)/\theta)$, $k \in \{1, \dots, m\}$. Note that this lower bound holds uniformly across the m regions. Hence, it is immediate to conclude that the same lower bound holds for the overall system. Since $\lim_{m \rightarrow +\infty} g(m) = 0$, it is clear that it is always possible to choose m so that $\phi_{\mathcal{B}(m)} \geq \phi^d$ (recall that $\mathbb{P}[P_j = 0] = 0$); in particular, a sufficient number of vehicles is given by the solution to the minimization problem in equation (7), and the theorem is proven. ■

The upper bound in equation (7) is valid under the least informative information structure, and *a fortiori* it is valid under any information structure. Hence, theorem 7.1 is valid under *any* information structure.

B. On the Constant $\bar{\beta}$ and the Use of Asymptotics

To compute \tilde{m} in equation (7), one needs to know, at least approximately, the value of $\bar{\beta}$; it is possible to show that when each region is approximately square-shaped, the value of $\bar{\beta}$ is approximately equal to $\sqrt{2}$ [9, page 765]. Furthermore, when λ is "large", one could reasonably use the asymptotic value $\beta_{\text{TSP}} \simeq 0.712$ (see Section II) to bound B_{r+1} ; it is then possible to show (the proof only requires minor modifications in the proof of theorem 7.1) that when λ is "large" one can replace $g(m)$ in equation (7) with $\tilde{g}(m) \doteq \beta_{\text{TSP}}^2 \lambda |\mathcal{E}| / (v^2 m^2)$.

Using theorem 7.1, we next show a scaling law for the minimum number of vehicles.

C. Scaling Law for the Minimum Number of Vehicles

Consider an information structure and a desired success factor ϕ^d ; moreover, let $\mathcal{I}(\lambda)$ be a problem instance where the arrival rate λ is a variable parameter, and let $\mathcal{P}(\lambda)$ be the corresponding set of admissible policies, parameterized by λ . The solution to \mathcal{OPT} is said to be $O(g(\lambda))$, where $g(\lambda) : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, if there exist $\Lambda \in \mathbb{R}_{> 0}$ and $c \in \mathbb{R}_{> 0}$ such that $\min_{\pi \in \mathcal{P}(\lambda)} m_{\pi}^*(\lambda) \leq c g(\lambda)$ for all $\lambda > \Lambda$. We have the following theorem.

Theorem 7.2 (Scaling law): When $s_{\max} = 0$ the solution to the optimization problem \mathcal{OPT} is $O(\sqrt{\lambda})$ under any information structure.

Proof: Define $\Theta \doteq \{\theta \in \mathbb{R}_{> 0} \mid 1 - F_P(\theta) > 0\}$; under the assumptions of the model, the set Θ is not empty; moreover, we have $0 < \sup \Theta < +\infty$. Let $\bar{\theta} = (1/2) \sup \Theta$. Then, we have $\sup_{\theta \in \mathbb{R}_{> 0}} (1 - F_P(\theta)) (1 - 2g(m; \lambda)/\theta) \geq (1 - F_P(\bar{\theta})) (1 - 2g(m; \lambda)/\bar{\theta}) \doteq h(m; \lambda)$, where we have made the dependency on λ explicit. Define (note that $1 - F_P(\bar{\theta}) > 0$) $m(\lambda) \doteq \lceil \sqrt{\lambda} \sqrt{|\mathcal{E}| \bar{\beta}} / (v \sqrt{\delta}) \rceil$, where $\delta = (1 - \eta \phi^d / (1 - F_P(\bar{\theta}))) \bar{\theta} / 2$, with $\eta > 1$. It is straightforward to show that there exists Λ such that for all $\lambda > \Lambda$ it holds $h(m(\lambda); \lambda) \geq \phi^d$. Hence we have, for $\lambda > \Lambda$, $\tilde{m}(\lambda) \leq m(\lambda)$, where $\tilde{m}(\lambda)$ is defined in theorem 7.1. Since $m_B^*(\lambda) \leq \tilde{m}(\lambda)$, we immediately obtain $\limsup_{\lambda \rightarrow +\infty} \frac{\min_{\pi \in \mathcal{P}(\lambda)} m_{\pi}^*(\lambda)}{\sqrt{\lambda}} \leq \sqrt{|\mathcal{E}| \bar{\beta}} / (v \sqrt{\delta})$. Since the batch policy is well-defined for any information structure, we conclude that when $s_{\max} = 0$ the solution to the optimization problem \mathcal{OPT} is $O(\sqrt{\lambda})$ under any information structure. ■

D. Simulations

We consider patience times that assume either the value 0.8 with 50% probability, or the value 1.6 with the remaining 50% probability; in other words, there are two types of demands, and one type is significantly more impatient than the other one. The arrival rate is $\lambda = 200$, the workspace is the unit square, the vehicles' velocity is $v = 1$, and $s_{\max} = 0$. Finally, we consider a desired success factor $\phi^d = 0.9$. By solving the minimization problem in equation (7) (with $\tilde{g}(m)$ instead of $g(m)$ since λ is "large"), we find $\tilde{m} = 36$. Figure 2 shows experimental values of ϕ_B as a function of the number of agents m . It can be noted that when $m = \tilde{m}$ the experimental success factor ϕ_B is larger than ϕ^d , in accordance with theorem 7.1. However, it is also possible to observe that the batch policy is able to guarantee a success factor larger than ϕ^d with a number of vehicles as low as 27; this is expected, since the techniques used in the proof of theorem 7.1 (e.g., Markov's inequality) lead to a conservative result.

VIII. CONCLUSION

We have studied a dynamic vehicle routing problem where demands have stochastic deadlines on their waiting times. This paper leaves numerous important extensions open for further research. First, in this paper we found a lower bound for the most informative case and we characterized two service policies that require the least amount of information;




Fig. 2. Experimental values of ϕ_B . The desired success factor is $\phi^d = 0.9$.

it would be very interesting to find lower bounds and study policies that are specific to each particular information structure. Second, our lower bound does not capture the dependency on λ and thus it is generally highly inaccurate for large values of the arrival rate; we are currently working on this issue. Third, we plan to remove some of the conservatism in the analysis of the batch policy. Finally, it is of interest to relax some of the assumptions in our model by considering, e.g., non-uniform distributions for demand locations and general renewal arrival processes. All these problems provide nontrivial challenges and might require techniques significantly different from those used in this paper.

REFERENCES

- [1] H. N. Psaraftis. Dynamic vehicle routing problems. In B. Golden and A. Assad, editors, *Vehicle Routing: Methods and Studies*, pages 223–248. Elsevier Science Publishers B.V., Amsterdam, The Netherlands, 1988.
- [2] D. J. Bertsimas and G. J. van Ryzin. Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles. *Operations Research*, 41(1):60–76, 1993.
- [3] M. Pavone, E. Frazzoli, and F. Bullo. Decentralized algorithms for stochastic and dynamic vehicle routing with general demand distribution. In *Proc. IEEE Conf. on Decision and Control*, pages 4869–4874, New Orleans, LA, December 2007.
- [4] P. Toth and D. Vigo. *The vehicle routing problem*. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, PA, 2002.
- [5] F. Baccelli, P. Boyer, and G. Hebuterne. Single-server queues with impatient customers. *Advances in Applied Probability*, 16(4):887–905, 1984.
- [6] A. R. Ward and P. W. Glynn. A diffusion approximation for a markovian queue with reneging. *Queueing Syst. Theory Appl.*, 43(1/2):103–128, 2003.
- [7] N. Bisnik, A. Abouzeid, and V. Isler. Stochastic event capture using mobile sensors subject to a quality metric. *IEEE Transactions on Robotics*, 23:676–692, 2007.
- [8] M. Pavone, N. Bisnik, E. Frazzoli, and V. Isler. A stochastic and dynamic vehicle routing problem with time windows and customer impatience. *Mobile Networks and Applications*, 14(3):350–364, 2009.
- [9] 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.
- [10] 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.
- [11] S. Asmussen. *Applied Probability and Queues*. Springer-Verlag, New York, NY, second edition, 2003.
- [12] S. Jr. Stidham. Regenerative processes in the theory of queues, with applications to the alternating-priority queue. *Advances in Applied Probability*, 4(3):542–577, 1972.
- [13] J. Cortés, S. Martínez, and F. Bullo. Spatially-distributed coverage optimization and control with limited-range interactions. *ESAIM. Control, Optimisation & Calculus of Variations*, 11:691–719, 2005.
- [14] N. N. Krasovskii. *Stability of Motion: Applications of Lyapunov's Second Method to Differential Systems and Equations with Delay*. Stanford University Press, CA, USA, 1963.
- [15] K. Laventall and J. Cortés. Coverage control by multi-robot networks with limited-range anisotropic sensory. *International Journal of Control*, 82(6):1113–1121, 2009.