

Adaptive and Distributed Algorithms for Vehicle Routing in a Stochastic and Dynamic Environment

Marco Pavone, Emilio Frazzoli, Francesco Bullo

Abstract

In this paper we present adaptive and distributed algorithms for motion coordination of a group of m vehicles. The vehicles must service demands whose time of arrival, spatial location and service requirement are stochastic; the objective is to minimize the average time demands spend in the system. The general problem is known as the m -vehicle Dynamic Traveling Repairman Problem (m -DTRP). The best previously known control algorithms rely on centralized task assignment and are not robust against changes in the environment. In this paper, we first devise new policies for the 1-DTRP that: (i) are provably optimal both in light load conditions (i.e., when the arrival rate for the demands is small) and in heavy load conditions (i.e., when the arrival rate for the demands is large), and (ii) are adaptive, in particular, they are robust against changes in load conditions. Then, we show that specific partitioning policies, whereby the environment is partitioned among the vehicles and each vehicle follows a certain set of rules in its own region, are optimal in heavy load conditions. Building upon the previous results, we finally design algorithms for the m -DTRP that (i) are adaptive and distributed, and (ii) have strong performance guarantees in heavy load conditions and stabilize the system in any load condition.

I. INTRODUCTION

Advances in computer technology and wireless communications have opened up tremendous possibilities for deploying large networks of autonomous mobile vehicles capable of interacting directly with the environment. Potential applications are numerous and include surveillance and exploration missions, where vehicles are required to visit points of interest, and transportation and distribution tasks, where vehicles are required to provide service or goods to people. In this context, one of the major scientific challenges is the design of possibly distributed coordination policies for the vehicles to assign and schedule demands for service (or *targets*) that are defined over an extended geographical area.

In the recent past, considerable efforts have been devoted to the problem of how to assign, cooperatively, demands to vehicles [1]–[6]. In these papers, the underlying mathematical model is *static* and fits within the framework of the Vehicle Routing Problem (see [7] for a thorough introduction to the VRP), whereby: (i) a team of m vehicles is required to service a set of n demands in a 2-dimensional space; (ii) each demand requires a certain amount of

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

Francesco Bullo is with the Center for Control Engineering and Computation, University of California at Santa Barbara, bullo@engineering.ucsb.edu

on-site service; (iii) the goal is to compute a set of routes that optimizes the cost of servicing (according to some quality of service metric) the demands. The VRP is *static* in the sense that vehicle routes are computed assuming that no new demands arrive.

Despite its richness, inherent elegance and frequent occurrence in practical problems, the VRP is often a *static approximation* to problems that are in reality *dynamic*, and therefore it fails to be an appropriate model in many applications of interest. In fact, demands often arrive *sequentially* in time, and planning algorithms should actually provide *policies* (in contrast to pre-planned routes) that prescribe how the routes should evolve as a function of those inputs that evolve in real-time. From a methodological point of view, dynamic demands (i.e., demands that vary over time) add *queueing phenomena* to the combinatorial nature of the VRP. As a motivating example, consider a surveillance mission where a team of Unmanned Aerial Vehicles (UAVs) must ensure continued coverage of a certain area; as events occur, i.e., as new targets are detected by on-board sensors or other assets (e.g., intelligence, high-altitude or orbiting platforms, etc.), UAVs must proceed to the location of the new event and provide close-range information about the target. Each request for close-range observation might also require an on-site service time (e.g., a time interval required to identify the target), which might be known a priori *only* through prior statistics. The UAV mission control is then a *continuous* process of detecting targets, planning routes and sending UAVs. In such a dynamic and stochastic setting, the routing process might rely on a priori statistics about the future stream of targets, and stability of the overall system is an additional issue to be considered.

Accordingly, in the present work we wish to address the assignment problem in the more general framework where new demands for service are generated *over time* by a *stochastic* process. In [8], the authors present a setting where demands arrive sequentially over time and routes are generated in real time, however the optimization problem is over a finite horizon and the analysis only addresses a number of vehicles $m \leq 2$. Arguably, the most general model for vehicle routing problems that have both a *dynamic* and a *stochastic* component is the *m*-vehicle Dynamic Traveling Repairman Problem (*m*-DTRP) [9]–[12]. In the *m*-DTRP, *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 (in the UAV example, the inter-arrival times might be exponential random variables, the demands' locations might be uniformly distributed within the workspace, and the time intervals required to identify a target might be exponential random variables as well). The objective is to find a policy to service demands over an infinite horizon that minimizes the expected system time (wait plus service) of the demands. The best previously known control policies for the *m*-DTRP rely on centralized task assignment and are not robust against changes in the environment, in particular changes in load conditions; therefore, they are of limited applicability in scenarios involving ad-hoc networks of autonomous vehicles operating in a time-varying environment.

In this paper we devise adaptive and distributed policies for the *m*-DTRP. We mainly focus on policies that in heavy load are constrained to provide *spatially-unbiased* service, i.e., the same quality of service to all outstanding demands independently of their locations. We argue in Section III how such constraint is indeed rather natural for applications. The contribution is threefold: First, we present a new class of policies for the 1-DTRP that are adaptive (e.g., to the load conditions) and satisfy the aforementioned constraint. In particular, we propose the Divide

TABLE I
ADAPTIVE POLICIES FOR THE 1-DTRP

Properties	DC Policy, $r \rightarrow \infty$	DC Policy, $r = 1$	RH Policy	UTSP Policy, $r \rightarrow \infty$
Light-load performance	optimal	optimal	optimal	not optimal
Heavy-load performance	optimal	within 100% of the optimal	within 100% of the optimal*	optimal
Adaptive	yes, except for f	yes	yes	no

& Conquer (DC) policy, whose performance depends on a design parameter $r \in \mathbb{N}_{>0}$. If $r \rightarrow +\infty$, the policy is (i) provably optimal both in light and heavy load conditions, and (ii) adaptive with respect to changes in the load conditions and in the statistics of the on-site service requirement; if, instead, $r = 1$, the policy is (i) provably optimal in light load conditions and within a factor 2 of the optimal in heavy load conditions, and (ii) adaptive with respect to all problem data, in particular, and perhaps surprisingly, it does not require *any* knowledge about the demand generation process. Moreover, by applying ideas of receding-horizon control to dynamic vehicle routing problems, we introduce the Receding Horizon (RH) policy, that also does not require any knowledge about the demand generation process; we show that the RH policy is optimal in light load and stable in any load condition, and we heuristically argue that its performance is close to optimality in heavy load conditions. Second, we show that specific partitioning policies, whereby the environment is partitioned among the vehicles and each vehicle follows a certain set of rules in its own region, are optimal in heavy load conditions. Finally, by combining the DC policy with suitable partitioning policies, we design a routing policy for the m -DTRP (called m -DC policy) that (i) is spatially distributed, scalable to large networks, and adaptive to network changes, (ii) is within a constant-factor of the optimal unbiased performance in heavy load conditions (in particular, it is optimal when demands are uniformly dispersed over the environment or when the average on-site service time \bar{s} is zero) and stabilizes the system in *any* load condition. Here, by network changes we mean changes in the number of vehicles, in the arrival rate of demands, and in the characterization of the on-site service requirement. We mention that the problems addressed in this paper, and the proposed algorithms, can not be cast within the framework of “traditional” queueing systems, and thus their analysis requires the use of non-standard techniques. Tables I and II provides a synoptic view about the properties of the main routing policies that we propose and analyze in this paper; in particular, our policies are compared with the best spatially-unbiased policy available in the literature, i.e., the UTSP policy with $r \rightarrow \infty$ [12]. (In Tables I and II an asterisk * signals that the result is heuristic, a dagger † means that the spatially-unbiased constraint might be violated, f denotes the spatial probability density for the demands’ locations, and m is the number of vehicles.)

The paper is structured as follows. In Section II we provide some background on probability theory, on the Euclidean Traveling Salesman Problem, and on the continuous multi-median problem. In Section III we formulate the problem we wish to address, and we review existing results. In Sections IV and V we present and analyze, respectively, the Divide & Conquer policy and the Receding Horizon policy. In Section VI, building upon the two

TABLE II
DISTRIBUTED AND ADAPTIVE POLICIES FOR THE m -DTRP

Properties	m -DC Policy, $r \rightarrow \infty$	m -RH Policy	UTSP Policy, $r \rightarrow \infty$
Light-load performance	“locally” optimal	“locally” optimal	not optimal
Heavy-load performance	optimal for uniform f , optimal for any f when $\bar{s} = 0$, within m of optimal unbiased in general [†]	optimal for uniform f^* , optimal for any f when $\bar{s} = 0^*$, within $2m$ of optimal unbiased in general ^{*,†}	optimal
Adaptive	yes, except for f	yes, except for f	no
Distributed	yes	yes	no

previous policies for the 1-DTRP, we discuss partitioning policies and we design spatially-distributed algorithms for the m -DTRP. In Section VII we discuss the particular case when the on-site service time \bar{s} is zero, and in Section VIII we present results from numerical experiments. Finally, in Section IX we draw some conclusions and discuss some directions for future work.

II. PRELIMINARIES

In this section, we briefly describe some known concepts from probability and locational optimization, on which we will rely extensively later in the paper.

A. Two Basic Results in Probability Theory

Suppose $h : \mathbb{R} \rightarrow \mathbb{R}$ is a *convex* function and X is a random variable, then *Jensen’s inequality* states that $\mathbb{E}[h(X)] \geq h(\mathbb{E}[X])$, provided both expectations exist. Moreover, if X is an integrable random variable (i.e., a random variable satisfying $\mathbb{E}[|X|] < +\infty$) and Y is any random variable, not necessarily integrable, on the same probability space, then $\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X|Y]]$.

B. Equitable Partitions

Let $\mathcal{Q} \subset \mathbb{R}^d$ be a bounded, convex set. An r -*partition* of \mathcal{Q} (where $r \in \mathbb{N}$, $r \geq 1$) is a collection of r closed subsets $\{\mathcal{Q}_k\}_{k=1}^r$ with disjoint interiors whose union is \mathcal{Q} . Given a measurable function $f : \mathcal{Q} \rightarrow \mathbb{R}_{>0}$, an r -partition $\{\mathcal{Q}_k\}_{k=1}^r$ is *equitable* with respect to f if $\int_{\mathcal{Q}_k} f(x) dx = \int_{\mathcal{Q}} f(x) dx / r$ for all $k \in \{1, \dots, r\}$. Similarly, given two measurable functions $f_j : \mathcal{Q} \rightarrow \mathbb{R}_{>0}$, $j \in \{1, 2\}$, an r -partition $\{\mathcal{Q}_k\}_{k=1}^r$ is *simultaneously equitable* with respect to f_1 and f_2 if $\int_{\mathcal{Q}_k} f_j(x) dx = \int_{\mathcal{Q}} f_j(x) dx / r$ for all $k \in \{1, \dots, r\}$ and $j \in \{1, 2\}$. Theorem 12 in [13] and Corollary 3 in [14] show that, given two measurable functions $f_j : \mathcal{Q} \rightarrow \mathbb{R}_{>0}$, $j \in \{1, 2\}$, there *always* exists an r -partition that is simultaneously equitable with respect to f_1 and f_2 and where the subsets \mathcal{Q}_k are convex.

C. Asymptotic and Worst-Case Properties of the Traveling Salesman Problem in the Euclidean Plane

The Euclidean Traveling Salesman Problem (for short, TSP) is formulated as follows: given a set D of n points in \mathbb{R}^d , find a minimum-length tour (i.e., a cycle that visits all nodes exactly once) of D ; the length of a tour is the

sum of all Euclidean distances on the tour. Let $\text{TSP}(D)$ denote the minimum length of a 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{Q} ; in [15] it is shown that there exists a constant $\beta_{\text{TSP},d}$ such that

$$\lim_{n \rightarrow +\infty} \frac{\mathbb{E}[\text{TSP}(D)]}{n^{1-1/d}} = \beta_{\text{TSP},d} \int_{\mathcal{Q}} \bar{f}(q)^{1-1/d} dq, \quad (1)$$

where \bar{f} is the density of the absolutely continuous part of the point distribution. For the case $d = 2$, the constant $\beta_{\text{TSP},2}$ has been estimated numerically as $\beta_{\text{TSP},2} \simeq 0.7120 \pm 0.0002$ [16]. Notice that the bound (1) holds for all compact sets: the shape of the set only affects the convergence rate to the limit. According to [17], if \mathcal{Q} is a “fairly compact and fairly convex” set in the plane, then equation (1) provides a “good” estimate of the average cost of the optimal TSP tour for values of n as low as 15.

Finally, for any bounded environment \mathcal{Q} , the following (deterministic) bound holds on the length of the TSP tour, uniformly on n [15]:

$$\text{TSP}(D) \leq \beta_{\mathcal{Q},d} |\mathcal{Q}|^{1/d} n^{1-1/d}, \quad (2)$$

where $|\mathcal{Q}|$ is the hypervolume of \mathcal{Q} and $\beta_{\mathcal{Q},d}$ is a constant generally larger than $\beta_{\text{TSP},d}$.

D. Tools for solving TSPs

The TSP is known to be NP-complete, 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. Even though the exact optimal solution of a large TSP can be very hard to compute, several exact and heuristic algorithms and software tools are available for the numerical solution of TSPs.

The most advanced TSP solver to date is arguably `concorde` [18]. Polynomial-time algorithms are available for constant-factor approximations of TSP solutions, among which we mention Christofides’ [19]. On a more theoretical side, Arora proved the existence of polynomial-time approximation schemes for the TSP, providing a $(1 + \varepsilon)$ constant-factor approximation for any $\varepsilon > 0$ [20].

A modified version of the Lin-Kernighan heuristic [21] is implemented in `linkern`; this powerful solver yields approximations in the order of 5% of the optimal tour cost very quickly for many instances. For example, in our numerical experiments on a machine with a 2GHz Intel Core Duo processor, approximations of random TSPs with 10,000 points typically required about twenty seconds of CPU time.¹

In the following, we will present algorithms that require on-line solutions of possibly large TSPs. Practical implementations of the algorithms will rely on heuristics or on polynomial-time approximation schemes, such as Lin-Kernighan’s or Christofides’. If a constant-factor approximation algorithm is used, the effect on the asymptotic performance guarantees of our algorithms can be simply modeled as a scaling of the constant $\beta_{\text{TSP},d}$.

¹Both `concorde` and `linkern` are freely available for academic research use at www.tsp.gatech.edu/concorde/index.html.

E. The Continuous Multi-Median Problem

Given a set $\mathcal{Q} \subset \mathbb{R}^d$ and a vector $P = (p_1, \dots, p_m)$ of m distinct points in \mathcal{Q} , the expected distance between a random point q , generated according to a probability density function f , and the closest point in P is given by

$$H_m(P, \mathcal{Q}) \doteq \mathbb{E} \left[\min_{k \in \{1, \dots, m\}} \|p_k - q\| \right] = \sum_{k=1}^m \int_{\mathcal{V}_k(P, \mathcal{Q})} \|p_k - q\| f(q) dq,$$

where $\mathcal{V}(P, \mathcal{Q}) = (\mathcal{V}_1(P, \mathcal{Q}), \dots, \mathcal{V}_m(P, \mathcal{Q}))$ is the Voronoi partition of the set \mathcal{Q} generated by the points P . In other words, $q \in \mathcal{V}_k(P, \mathcal{Q})$ if $\|q - p_k\| \leq \|q - p_j\|$, for all $j \in \{1, \dots, m\}$. The set \mathcal{V}_k is referred to as the Voronoi cell of the generator p_k . The function H_m is known in the locational optimization literature as the continuous Weber function or the continuous multi-median function; see [22], [23] and references therein.

The m -median of the set \mathcal{Q} , with respect to the measure induced by f , is the global minimizer

$$P_m^*(\mathcal{Q}) = \arg \min_{P \in \mathcal{Q}^m} H_m(P, \mathcal{Q}).$$

We let $H_m^*(\mathcal{Q}) = H_m(P_m^*(\mathcal{Q}), \mathcal{Q})$ be the global minimum of H_m . It is straightforward to show that the map $P \mapsto H_1(P, \mathcal{Q})$ is differentiable and strictly convex on \mathcal{Q} . Therefore, it is a simple computational task to compute $P_1^*(\mathcal{Q})$. It is convenient to refer to $P_1^*(\mathcal{Q})$ as the median of \mathcal{Q} . On the other hand, when $m > 1$, the map $P \mapsto H_m(P, \mathcal{Q})$ is differentiable (whenever (p_1, \dots, p_m) are distinct) but not convex, thus making the solution of the continuous m -median problem hard in the general case. It is known [22], [24] that the discrete version of the m -median problem is NP-hard for $d \geq 2$. The set of critical points of H_m contains all configurations (p_1, \dots, p_m) with the property that each point p_k is the generator of the Voronoi cell $\mathcal{V}_k(P, \mathcal{Q})$ as well as the median of $\mathcal{V}_k(P, \mathcal{Q})$ (we refer to such Voronoi diagrams as *median* Voronoi diagrams).

III. PROBLEM FORMULATION AND EXISTING RESULTS

In this section we formulate the problem we wish to address and we review existing results.

A. Problem Formulation

The basic version of the problem we wish to study in this paper is known as the m -vehicle Dynamic Traveling Repairman Problem (m -DTRP), and was introduced by Psaraftis in [9] and mainly studied by Bertsimas and van Ryzin in [10]–[12].

Let the workspace $\mathcal{Q} \subset \mathbb{R}^d$ be a bounded, convex set, and let $\|\cdot\|$ denote the Euclidean norm in \mathbb{R}^d . We define the diameter of \mathcal{Q} as: $\text{diam}(\mathcal{Q}) \doteq \sup\{\|p - q\| \mid p, q \in \mathcal{Q}\}$. For simplicity, we will only consider the planar case, i.e., $d = 2$, with the understanding that extensions to higher dimensions are possible. Consider m holonomic vehicles, modeled as point masses, and let

$$p(t) = (p^1(t), \dots, p^m(t)) \in \mathcal{Q}^m$$

describe the locations of the vehicles at time t . The vehicles are free to move, traveling at a constant velocity v , within the environment \mathcal{Q} . The vehicles are identical, and have unlimited range and demand servicing capacity.

Demands are generated according to a homogeneous (i.e., time-invariant) spatio-temporal Poisson process², with time intensity $\lambda \in \mathbb{R}_{>0}$, and spatial density $f : \mathcal{Q} \rightarrow \mathbb{R}_{>0}$. In other words, demands arrive to \mathcal{Q} according to a Poisson process with intensity λ , and their locations $\{X_j; j \geq 1\}$ are i.i.d. (i.e., *independent and identically distributed*) and distributed according to a density f whose support is \mathcal{Q} . A demand's location becomes known (is realized) at its arrival epoch; thus, at time t we know with *certainty* the locations of demands that arrived prior to time t , but future demand locations form an i.i.d sequence. The notation $f(x)$ is short for $f(x_1, x_2)$ ($x = [x_1, x_2]$). The density f satisfies:

$$\mathbb{P}[X_j \in S] = \int_S f(x) dx \quad \forall S \subseteq \mathcal{Q}, \quad \text{and} \quad \int_{\mathcal{Q}} f(x) dx = 1.$$

We assume that the number of initial outstanding demands is a random variable with finite first and second moments.

At each demand location, vehicles spend some time s in on-site service that is i.i.d. and generally distributed with finite first and second moments denoted by \bar{s} and \bar{s}^2 (we also assume $\bar{s} > 0$; we will discuss in Section VII the case $\bar{s} = 0$). A realized demand is removed from the system after one of the vehicles has completed its on-site service. 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. 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 $\rho \rightarrow 0^+$, in the sense that $\lambda \rightarrow 0^+$; when we refer to *heavy load* conditions, we consider the case $\rho \rightarrow 1^-$, in the sense that $\lambda \rightarrow [m/\bar{s}]^-$.

Let \mathcal{P} be the set of all causal, stable and stationary routing policies. Letting \bar{T}_π , denote the system time of a particular policy $\pi \in \mathcal{P}$, the m -DTRP is then defined as the problem of finding 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 in the right hand side above.

We finally need two definitions; in these definitions, X is the location of a randomly chosen demand and W is its waiting time.

Definition 3.1: A policy π is called *spatially unbiased* if for every pair of sets $\mathcal{S}_1, \mathcal{S}_2 \subseteq \mathcal{Q}$

$$\mathbb{E}[W | X \in \mathcal{S}_1] = \mathbb{E}[W | X \in \mathcal{S}_2];$$

²There are three main reasons behind the choice of modeling the demand arrival process with a Poisson process: First, the system we deal with is a spatial queueing system, and the arrival process to a queueing systems is often well approximated by a Poisson process (see, e.g., [25]). Second, the Poisson process leads to a model that is analytically tractable. Third, the Poisson process, being “the most random” way to distribute events in time, leads to useful “worst-case” scenario evaluations (worst-case with respect to the possible types of arrival processes). As a final consideration, most of the results in this paper extend with little effort to the more general class of renewal processes.

a policy π is instead called *spatially biased* if there exists sets $\mathcal{S}_1, \mathcal{S}_2 \subseteq \mathcal{Q}$ such that

$$\mathbb{E}[W | X \in \mathcal{S}_1] > \mathbb{E}[W | X \in \mathcal{S}_2].$$

B. Existing Results for the m -DTRP

As in many queueing problems, the analysis of the m -DTRP problem for all values of $\rho \in (0, 1)$ is difficult. In [10]–[12], [26], lower bounds for the optimal steady-state system time are derived for the light load case (i.e., $\rho \rightarrow 0^+$), and for the heavy load case (i.e., $\rho \rightarrow 1^-$). Subsequently, policies are designed for these two limiting regimes, and their performance is compared to the lower bounds.

1) *Lower Bounds:* For the light load case, a tight lower bound on the system time is derived in [11]. In the light load case, the lower bound on the system time is strongly related to the solution of the m -median problem:

$$\bar{T}^* \geq \frac{1}{v} H_m^*(\mathcal{Q}) + \bar{s}, \quad \text{as } \rho \rightarrow 0^+. \quad (3)$$

This bound is tight: there exist policies whose system times, in the limit $\rho \rightarrow 0^+$, attain this bound; we present such asymptotically optimal policies for the light load case below.

The following two lower bounds hold in the heavy load case [12], [26]. Within the class in \mathcal{P} of spatially *unbiased* policies the optimal system time is lower bounded by

$$\bar{T}_U^* \geq \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda \left[\int_{\mathcal{Q}} f^{1/2}(x) dx \right]^2}{m^2 v^2 (1 - \rho)^2} \quad \text{as } \rho \rightarrow 1^-. \quad (4)$$

Within the class in \mathcal{P} of spatially *biased* policies the optimal system time is lower bounded by

$$\bar{T}_B^* \geq \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda \left[\int_{\mathcal{Q}} f^{2/3}(x) dx \right]^3}{m^2 v^2 (1 - \rho)^2} \quad \text{as } \rho \rightarrow 1^-. \quad (5)$$

Both bounds (4) and (5) are tight: there exist policies whose system times, in the limit $\rho \rightarrow 1^-$, attain these bounds; therefore the inequalities in (4) and (5) should indeed be replaced by equalities. We present asymptotically optimal policies for the heavy load case below. It is shown in [12] that the lower bound in equation (5) is always lower than or equal to the lower bound in equation (4) for all densities f .

Remark 3.2: In equations (4) and (5), the right-hand side approaches $+\infty$ as $\rho \rightarrow 1^-$. Thus, one should more formally write the inequalities with $\bar{T}^*(1 - \rho)^2$ on the left-hand side, so that the right-hand side is finite. However, this makes the presentation less readable, and thus, in the remainder of the paper, we adhere to the less formal but more transparent style of equations (4) and (5). Furthermore, this is the style generally used in the literature on the m -DTRP (see, e.g., [12] Section 4 and [27] Proposition 2).

Remark 3.3: It is possible to show (see [12], Proposition 1) that a *uniform* spatial density function is the *worst possible* and that any deviation from uniformity in the demand distribution will strictly lower the optimal mean system time in both the unbiased and biased case. In addition, allowing biased service will result in a strict reduction of the optimal mean system time for any non-uniform density f . Also, when the density is uniform there is nothing to be gained by providing biased service.

Remark 3.4: In this paper we assume that the vehicles have full knowledge about the spatial demand distribution f . In practice, the vehicles might use an empirical distribution function for the demand locations (which might be computed either through a centralized algorithm or through a distributed algorithm). Since the demand locations are i.i.d., the empirical distribution will converge to the true distribution, and therefore after a transient (whose length depends on the arrival rate λ) the routing policies described in this paper will behave according to the bounds derived under the assumption that f is known by the vehicles.

2) *Optimal Policies for the Light Load Case:* For the light load case, to the best of our knowledge, there is only one policy that is known to be optimal; this policy is as follows.

The m Stochastic Queue Median (SQM) Policy [11] — Locate one vehicle at each of the m median locations for the environment \mathcal{Q} . When demands arrive, assign them to the nearest median location and its corresponding vehicle. Have each vehicle service its respective demands in First-Come, First-Served (FCFS) order returning to its median after each service is completed.

This policy, although optimal in light load, has two characteristics that limit its application to large-scale vehicle networks: First, it quickly becomes unstable as the load increases, i.e., there exists $\rho_c < 1$ such that for all $\rho > \rho_c$ the system time \bar{T}_{SQM} is infinite (hence, this policy is not adaptive with respect to the load conditions). Second, a *central entity* needs to compute the m median locations and assign them to the vehicles (hence, it is centralized).

3) *Optimal Policies for the Heavy Load Case:* For the heavy load case, to the best of our knowledge, there is only one unbiased policy that has been proven to achieve the lower bound in (4); this policy is as follows.

The Unbiased TSP (UTSP) Policy [12] — Let r be a fixed positive, large integer. From a central point in the interior of \mathcal{Q} , subdivide the service region into r wedges $\mathcal{Q}_1, \mathcal{Q}_2, \dots, \mathcal{Q}_r$ such that $\int_{\mathcal{Q}_k} f(x)dx = 1/r$, $k \in \{1, \dots, r\}$. Within each subregion, form sets of size n/r (n is a design parameter). As sets are formed, deposit them in a queue and service them FCFS with the first available vehicle by forming a TSP on the set and following it in an arbitrary direction. Optimize over n .

It is possible to show that

$$\bar{T}_{\text{UTSP}} \leq \left(1 + \frac{m}{r}\right) \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda \left[\int_{\mathcal{Q}} f^{1/2}(x)dx\right]^2}{m^2 v^2 (1 - \rho)^2} \quad \text{as} \quad \rho \rightarrow 1^-.$$

Thus, letting $r \rightarrow \infty$, the lower bound in (4) is achieved.

The same paper [12] presents an optimal biased policy. This policy, called Biased TSP (BTSP) Policy, relies on an even finer partition of the environment, requires f to be piecewise constant, and also depends on a parameter n .

Although both the UTSP and the BTSP policies are optimal within their respective classes, they have two major drawbacks: First, in the UTSP policy, to ensure stability, n should be chosen so that (see [12], page 961)

$$n > \frac{\lambda^2 \beta_{\text{TSP},2}^2 \left[\int_{\mathcal{Q}} f^{1/2}(x) dx\right]^2}{m^2 v^2 (1 - \rho)^2};$$

therefore, to ensure stability over a wide range of values of ρ , the system designer is forced to select a *large* value for n . However, if during the execution of the policy the load factor turns out to be only moderate, demands have

to wait for an excessively large set to be formed, and the overall system performance deteriorates significantly. Similar considerations hold for the BTSP policy. Hence, these two policies are not adaptive with respect to the load conditions. Second, both policies require a centralized data structure (the demands' queue is shared by the vehicles); hence, both policies are centralized.

C. Objective of the Paper

In many applications it is desirable to provide the same quality of service to all of the outstanding demands independently of their locations; e.g., in the UAV example discussed in the introduction, outstanding “threats” at the periphery of \mathcal{Q} should receive the same quality of service as the outstanding “threats” in the middle of \mathcal{Q} . However, in light load, when most of the time the vehicles are idling, it is more reasonable to consider *spatially-biased* policies that aim at positioning the vehicles in the best a-priori locations, i.e., at the locations that minimize the expected distance to the next demand.

Accordingly, in this paper we focus on policies for the DTRP that are constrained to provide spatially-unbiased service *in heavy load*, but can provide spatially-biased service in light load. With a slight abuse, we refer to policies satisfying this constraint as *unbiased policies*. The objective is to devise *adaptive, spatially-distributed, scalable* unbiased policies with *provable performance guarantees*, that rely only on local exchange of information between neighboring vehicles, and do not require explicit knowledge of the global structure of the network.

D. Toward Adaptive, Distributed, Scalable Control Policies for the m -DTRP

A candidate adaptive, distributed, scalable control policy for the m -DTRP is the simple Nearest Neighbor (NN) policy: at each service completion epoch, each vehicle chooses to visit next the closest unserved demand, if any, otherwise it stops at the current position. Because of the dependencies among the travel inter-demand distances, the analysis of the NN policy is difficult and no rigorous results have been obtained so far [10]; in particular, there are no rigorous results about its stability properties. Simulation experiments show that the NN policy performs like a biased policy in any load, and is not optimal *neither* in the light load case *nor* in the heavy load case. In particular, when f is uniform, in light load the steady-state system time under the NN policy is 40% larger than the steady-state system time under the (optimal) SQM policy, while in heavy load the steady-state system time under the NN policy is 50% larger than the steady-state system time under the (optimal) Biased TSP policy [10], [12]. Therefore, the NN policy *lacks* provable performance guarantees (in particular about stability), is *biased* and does *not* seem to achieve optimal performance neither in light load nor in heavy load.

The key idea that we will pursue in this paper is that of *partitioning policies*.

Definition 3.5: Given a policy π for the 1-DTRP and m vehicles, a π -partitioning policy is a family of multi-vehicle policies such that

- 1) the environment \mathcal{Q} is partitioned into m openly disjoint subregions \mathcal{Q}_k , $k \in \{1, \dots, m\}$, whose union is \mathcal{Q} ,
- 2) one vehicle is assigned to each subregion (thus, there is a one-to-one correspondence between vehicles and subregions),

3) each vehicle executes the single-vehicle policy π to service demands that fall within its own subregion.

Note that Definition 3.5 does not specify how the environment is actually partitioned, therefore Definition 3.5 describes a *family* of policies (one for each partitioning strategy) for the m -DTRP. The SQM policy, which is optimal *in light load*, is indeed a partitioning policy whereby \mathcal{Q} is partitioned according to a median Voronoi diagram that “minimizes” H_m , and each vehicle executes inside its own Voronoi cell the single-vehicle policy “service FCFS and return to the median after each service completion”. Moreover, in Section VI, we will prove that *in heavy load*, given a single-vehicle, unbiased optimal policy π^* and m vehicles, there exists an unbiased π^* -partitioning policy that is optimal.

The above discussion leads to the following strategy: First, for the 1-DTRP, we design unbiased and adaptive control policies with provable performance guarantees. Then, by using recently developed *spatially-distributed* algorithms for environment partitioning, we extend these single-vehicle policies to spatially-distributed multi-vehicle policies.

IV. THE SINGLE-VEHICLE DIVIDE & CONQUER POLICY

We describe the Divide & Conquer (DC) policy in Algorithm 1, where D is the set of outstanding demands waiting for service and p is the current position of the vehicle. An r -partition that is simultaneously equitable

Algorithm 1: Divide & Conquer (DC) Policy

Assumes: An r -partition $\{\mathcal{Q}_k\}_{k=1}^r$ of \mathcal{Q} that is simultaneously equitable with respect to f and $f^{1/2}$.

```

1 if  $D = \emptyset$  then
2   | Let  $\tilde{P}_1^*$  be the point minimizing the sum of distances to demands serviced in the past (the initial condition
   | for  $\tilde{P}_1^*$  is a random point in  $\mathcal{Q}$ ); if  $p \neq \tilde{P}_1^*$ , move to  $\tilde{P}_1^*$ , otherwise stop.
3 else
4   |  $k \leftarrow k_0$ . //  $k_0$  is uniformly randomly chosen in the set  $\{1, \dots, r\}$ .
5   repeat
6     | Move to the median of subregion  $\mathcal{Q}_k$ .
7     | Compute a TSP tour through all demands in subregion  $\mathcal{Q}_k$ . Service all demands by following the TSP
   |   | tour, starting at a random point on the TSP tour.
8     |  $k \leftarrow k + 1$  modulo  $r$ .
9   until  $D = \emptyset$ 
10 Repeat.
```

with respect to f and $f^{1/2}$ is indeed guaranteed to exist (see discussion on equitable partitions in Section II). The number of subregions r is a design parameter whose choice will be discussed after the analysis of the policy. One should note that \tilde{P}_1^* (which is defined in Algorithm 1) is indeed the empirical median. Next, we characterize the DC policy. Notice that the DC policy is an *unbiased* policy.

A. Analysis of the DC Policy in Light Load

We first study the light load case (i.e., $\rho \rightarrow 0^+$).

Theorem 4.1 (Performance of DC policy in light load): As $\rho \rightarrow 0^+$, for every r , the DC policy is asymptotically optimal, that is,

$$\bar{T}_{\text{DC}}(r) \rightarrow \bar{T}^*, \quad \text{as } \rho \rightarrow 0^+.$$

Proof: Using arguments similar to those in [27], consider generic initial conditions in \mathcal{Q} for the vehicle's position and for the positions of the outstanding demand (denote the initial number of demands with n_0). Under the DC policy, an upper bound to the time C_0 needed to service *all* of the initial demands and move to the empirical median \tilde{P}_1^* is: $C_0 \leq (n_0 + 1) \text{diam}(\mathcal{Q})/v + \sum_{j=1}^{n_0} s_j$. Thus, we have $\mathbb{E}[C_0] \leq \mathbb{E}[n_0](\text{diam}(\mathcal{Q})/v + \bar{s}) + \text{diam}(\mathcal{Q})/v$.

Given the time interval $C_0 = t$, $t \in \mathbb{R}_{\geq 0}$, the probability that no *new* demand arrives during such time interval is $\mathbb{P}(N(C_0) = 0 | C_0 = t) = \exp(-\lambda t)$, where $N(t)$ is the counting variable associated to a Poisson process with rate λ . Hence, by applying the law of total probability and Jensen's inequality for convex functions

$$\begin{aligned} \mathbb{P}[N(C_0) = 0] &= \int_0^\infty \mathbb{P}(N(C_0) = 0 | C_0 = t) f_{C_0}(t) dt = \int_0^\infty \exp(-\lambda t) f_{C_0}(t) dt \geq \exp\left(-\lambda \int_0^\infty t f_{C_0}(t) dt\right) \\ &= \exp\left(-\lambda \mathbb{E}[C_0]\right) \geq \exp\left(-\lambda \mathbb{E}[n_0](\text{diam}(\mathcal{Q})/v + \bar{s}) - \lambda \text{diam}(\mathcal{Q})/v\right), \end{aligned} \quad (6)$$

where f_{C_0} is the density function of random variable C_0 . Since, by assumption, $\mathbb{E}[n_0] < \infty$ and $\rho \rightarrow 0^+$ (i.e., $\lambda \rightarrow 0^+$), we obtain $\mathbb{P}[N(C_0) = 0] \rightarrow 1^-$ as $\rho \rightarrow 0^+$. As a consequence, after an initial transient, all demands will be generated with the vehicle at the empirical median \tilde{P}_1^* and with an empty demand queue, with probability 1. Then, in light load, the DC policy becomes identical to the light load routing policy presented in [28]; since such routing policy is optimal, in light load, for the 1-DTRP [28], we obtain the desired result (the proof in [28] basically shows that $\tilde{P}_1^* \rightarrow P_1^*(\mathcal{Q})$). ■

B. Analysis of the DC Policy in Heavy Load

Next, we consider the DC policy in heavy load, i.e., when $\rho \rightarrow 1^-$. The performance of the DC policy in heavy load is characterized by the following theorem.

Theorem 4.2 (Performance of DC policy in heavy load): As $\rho \rightarrow 1^-$, the system time for the DC policy satisfies

$$\bar{T}_{\text{DC}}(r) \leq \left(1 + \frac{1}{r}\right) \frac{\beta_{\text{TSP},2}^2 \lambda \left[\int_{\mathcal{Q}} f^{1/2}(x) dx\right]^2}{2 v^2 (1 - \rho)^2}, \quad (7)$$

where r is the number of subregions.

Contrary to the UTSP policy, the DC policy does not define a $\sum GI/G/1$ queue (a queue is denoted $\sum GI/G/1$ if its input process is the superposition of r independent renewal processes). Therefore, to prove Theorem 4.2 we use non-standard techniques similar to those we developed in [29], [30]. The proof of Theorem 4.2 builds on a number of intermediate results; we start with the following lemma, similar to Lemma 1 in [27], characterizing the number of outstanding demands in heavy load.

Lemma 4.3 (Number of outstanding demands in heavy load for DC policy): In heavy load (i.e., $\rho \rightarrow 1^-$), after a transient, the number of demands serviced in a single tour of the vehicle in the DC policy 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: Consider the case where the vehicle moves with infinite velocity (i.e., $v \rightarrow +\infty$); then the system is reduced to the usual M/G/1 queue. The infinite-velocity system has fewer demands waiting in queue. A known result on M/G/1 queues [31] states that, after transients, the total number of demands, as $\rho \rightarrow 1^-$, is very large with high probability. Thus, the number of outstanding demands in each subregion is also very large with high probability. In particular, after transients, the number of demands is very large with high probability at the time instants when the vehicle starts a new TSP tour. ■

Lemma 4.3 has two implications. First, since the number of demands is very large at the time instants when the vehicle starts a new TSP tour, we can apply equation (1) to estimate the average length of a TSP tour. Second, since $D \neq \emptyset$ with high probability, the DC policy, in heavy load, *reduces* to lines 5-9 of Algorithm 1 with the condition $D = \emptyset$ always false.

We refer to the time instant t_i , $i \geq 0$, in which the vehicle starts a new TSP tour in subregion 1 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^k , $k \in \{1, \dots, r\}$, be the number of outstanding demands serviced in subregion k during iteration i . Finally, let C_i^k be the time interval between the time instant the vehicle starts to service demands in subregion k during iteration i and the time instant the vehicle starts to service demands in the same subregion k during next iteration $i+1$. Demands arrive in subregion \mathcal{Q}_k according to a Poisson process with rate $\hat{\lambda} \doteq \lambda/r$, where we use the fact that the partition $\{\mathcal{Q}_k\}_{k=1}^r$ is equitable with respect to f ; then, we have $\mathbb{E}[n_{i+1}^k] = \hat{\lambda} \mathbb{E}[C_i^k]$. The time interval C_i^k is the sum of three components:

- 1) the time for the r travels from the end of one TSP tour to the start of next one;
- 2) the time to perform r TSP tours;
- 3) the time to provide on-site service while performing the r TSP tours.

The first component is trivially upper bounded by $r \cdot 2 \text{diam}(\mathcal{Q})/v$. Given that a demand falls in subregion \mathcal{Q}_k , the conditional density for its location (whose support is \mathcal{Q}_k) is $f(x)/\left(\int_{\mathcal{Q}_k} f(x) dx\right)$. By equation (1), we have that the expected length of a TSP tour through n demands distributed according to the density $f(x)/\left(\int_{\mathcal{Q}_k} f(x) dx\right)$ satisfies

$$\lim_{n \rightarrow +\infty} \frac{\mathbb{E}[\text{TSP}(n)]}{\sqrt{n}} = \beta_{\text{TSP},2} \left(\int_{\mathcal{Q}_k} \sqrt{\frac{f(x)}{\int_{\mathcal{Q}_k} f(x) dx}} dx \right) = \beta_{\text{TSP},2} \sqrt{r} \frac{\int_{\mathcal{Q}} f^{1/2}(x) dx}{r} \doteq \beta, \quad (8)$$

where we have exploited the fact that, by definition of the DC policy, the partition $\{\mathcal{Q}_k\}_{k=1}^r$ is simultaneously equitable with respect to f and $f^{1/2}$. Pick now an arbitrarily small $\zeta > 0$. By equation (8), there exists an $\tilde{n} \in \mathbb{N}_{>0}$ such that for all $n > \tilde{n}$ it holds $\mathbb{E}[\text{TSP}(n)] \leq (\beta + \zeta) \sqrt{n}$. Then, the expected length of a TSP tour through n_i^k

demands (with a slight abuse of notation, we call such length $\text{TSP}(n_i^k)$) can be upper bounded as

$$\begin{aligned} \mathbb{E} [\text{TSP}(n_i^k)] &= \sum_{n=0}^{+\infty} \mathbb{E} [\text{TSP}(n_i^k) | n_i^k = n] \mathbb{P} [n_i^k = n] \leq \sum_{n=\tilde{n}+1}^{+\infty} (\beta + \zeta) \sqrt{n} \mathbb{P} [n_i^k = n] + \sum_{n=0}^{\tilde{n}} \tilde{n} \text{diam}(\mathcal{Q}) \mathbb{P} [n_i^k = n] \\ &\leq \sum_{n=0}^{+\infty} (\beta + \zeta) \sqrt{n} \mathbb{P} [n_i^k = n] + \tilde{n} \text{diam}(\mathcal{Q}) \mathbb{P} [n_i^k \leq \tilde{n}] = (\beta + \zeta) \mathbb{E} [\sqrt{n_i^k}] + \tilde{n} \text{diam}(\mathcal{Q}) \mathbb{P} [n_i^k \leq \tilde{n}] \\ &\leq (\beta + \zeta) \sqrt{\mathbb{E} [n_i^k]} + \tilde{n} \text{diam}(\mathcal{Q}) \mathbb{P} [n_i^k \leq \tilde{n}], \end{aligned} \tag{9}$$

where in the first equality we have used the law of iterated expectation, and in the last inequality we have applied Jensen's inequality for concave functions in the form $\mathbb{E} [\sqrt{X}] \leq \sqrt{\mathbb{E} [X]}$. By Lemma 4.3, there exists $\bar{\varrho} \in (0, 1)$ such that for all $\varrho \in [\bar{\varrho}, 1)$ there exists an integer $\bar{i}(\varrho)$ (possibly dependent on ϱ) such that $\mathbb{P} [n_i^k \leq \tilde{n}] < 1/\tilde{n}$. Consider any $\varrho \in [\bar{\varrho}, 1)$ and assume that $i \geq \bar{i}(\varrho)$, then

$$\mathbb{E} [\text{TSP}(n_i^k)] \leq (\beta + \zeta) \sqrt{\mathbb{E} [n_i^k]} + \text{diam}(\mathcal{Q}), \quad \text{for } i \geq \bar{i}(\varrho) \text{ and } \varrho \in [\bar{\varrho}, 1).$$

Finally, since on-site service times are independent of the number of outstanding demands, the expected on-site service requirement for n_i^k demands is $\bar{s} \mathbb{E} [n_i^k]$.

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

$$\bar{n}_{i+1}^k = \hat{\lambda} \mathbb{E} [C_i^k] \leq \hat{\lambda} \left(\text{diam}(\mathcal{Q})(2r/v + 1) + \frac{\beta + \zeta}{v} \sum_{j=k}^r \sqrt{\bar{n}_i^j} + \frac{\beta + \zeta}{v} \sum_{j=1}^{k-1} \sqrt{\bar{n}_{i+1}^j} + \bar{s} \sum_{j=k}^r \bar{n}_i^j + \bar{s} \sum_{j=1}^{k-1} \bar{n}_{i+1}^j \right), \tag{10}$$

where $k \in \{1, \dots, r\}$, $i \geq \bar{i}(\varrho)$, and $\varrho \in [\bar{\varrho}, 1)$.

The r inequalities above describe a system of recurrence relations that allows us to find an upper bound on $\limsup_{i \rightarrow +\infty} \bar{n}_i^k$. The following lemma bounds the values to which the limits $\limsup_{i \rightarrow +\infty} \bar{n}_i^k$ converge.

Lemma 4.4 (Steady state number of demands for DC policy): In heavy load, for every set of initial conditions $\{\bar{n}_1^k\}_{k \in \{1, \dots, r\}}$, the trajectories $i \mapsto \bar{n}_i^k$, $k \in \{1, \dots, r\}$ satisfy

$$\bar{n}^k \doteq \limsup_{i \rightarrow +\infty} \bar{n}_i^k \leq \beta_{\text{TSP}, 2}^2 \frac{\lambda^2 \left[\int_{\mathcal{Q}} f^{1/2}(x) dx \right]^2}{r v^2 (1 - \varrho)^2}.$$

Proof: Define the constants: $\tilde{\beta} \doteq \beta + \zeta$ and $a \doteq \text{diam}(\mathcal{Q})(2r/v + 1)$. Henceforth we assume that $i \geq \bar{i}(\varrho)$ and that $\varrho \in [\bar{\varrho}, 1)$, and we simply denote $\bar{i}(\varrho)$ as \bar{i} . Next we define two auxiliary systems, System-Y and System-Z. The initial conditions of these two systems will be set equal to \bar{n}_1^k , and we will use their trajectories to bound the trajectories $i \mapsto \bar{n}_i^k$. We define System-Y (with state $y(i) \in \mathbb{R}^r$) as

$$y_k(i+1) = \hat{\lambda} a + \hat{\lambda} \left(\sum_{j=k}^r \left(\bar{s} y_j(i) + \frac{\tilde{\beta}}{v} \sqrt{y_j(i)} \right) + \sum_{j=1}^{k-1} \left(\bar{s} y_j(i+1) + \frac{\tilde{\beta}}{v} \sqrt{y_j(i+1)} \right) \right), \quad k \in \{1, \dots, r\}. \tag{11}$$

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

$$\sqrt{a} \leq \frac{1}{4\varepsilon} + \varepsilon a, \quad \text{for all } a \in \mathbb{R}_{\geq 0}.$$

By applying Young's inequality in equation (11) we obtain

$$y_k(i+1) \leq \hat{\lambda} \left(\bar{s} + \varepsilon \frac{\tilde{\beta}}{v} \right) \left(\sum_{j=k}^r y_j(i) + \sum_{j=1}^{k-1} y_j(i+1) \right) + \hat{\lambda} \left(\frac{r \tilde{\beta}}{4v\varepsilon} + a \right), \quad k \in \{1, \dots, r\}. \quad (12)$$

Next, define System-Z (with state $z(i) \in \mathbb{R}^r$) as

$$z_k(i+1) = \hat{\lambda} \left(\bar{s} + \varepsilon \frac{\tilde{\beta}}{v} \right) \left(\sum_{j=k}^r z_j(i) + \sum_{j=1}^{k-1} z_j(i+1) \right) + \hat{\lambda} \left(\frac{r \tilde{\beta}}{4v\varepsilon} + a \right), \quad k \in \{1, \dots, r\}. \quad (13)$$

The proof now proceeds as follows. First, we show that if $\bar{n}_i^k \leq y_k(i) \leq z_k(i)$ for all k , then

$$\bar{n}_i^k \leq y_k(i) \leq z_k(i), \quad \text{for all } k \text{ and for all } i \geq \bar{i}. \quad (14)$$

Second, we show that the trajectories of System-Z are bounded; this fact, together with equation (14), implies that also trajectories of variables \bar{n}_i^k and $y_k(i)$ are bounded. Third, and last, we will compute $\limsup_{i \rightarrow +\infty} y_k(i)$; this quantity, together with equation (14), will yield the desired result.

Let us consider the first issue. Given i , we prove that if $\bar{n}_i^k \leq y_k(i) \leq z_k(i)$ for all $k \in \{1, \dots, r\}$, then $\bar{n}_{i+1}^k \leq y_k(i+1) \leq z_k(i+1)$. Indeed, it is immediate to show that, under the assumptions, it holds: $\bar{n}_{i+1}^1 \leq y_1(i+1) \leq z_1(i+1)$; moreover, it is easy to show that, under the assumptions, and if $\bar{n}_{i+1}^j \leq y_j(i+1) \leq z_j(i+1)$ for all $j \leq k < r$, then $\bar{n}_{i+1}^{k+1} \leq y_{k+1}(i+1) \leq z_{k+1}(i+1)$. The claim follows by induction on k . Equation (14) is a consequence of the statement we have just proven.

We now turn our attention to the second issue, namely boundedness of trajectories for System-Z (described in equation (13)). Define $\delta \doteq \hat{\lambda}(\bar{s} + \varepsilon \tilde{\beta}/v)$. System-Z is a discrete-time linear system and can be rewritten (by adding and subtracting $z_{k-1}(i)$ in the second term between parentheses in equation (13), when $k > 1$) as

$$z_k(i+1) = \begin{cases} \delta \sum_{j=1}^r z_j(i) + \hat{\lambda} \left(\frac{r \tilde{\beta}}{4v\varepsilon} + a \right), & \text{for } k = 1, \\ (1 + \delta) z_{k-1}(i+1) - \delta z_{k-1}(i), & \text{for } k \in \{2, \dots, r\}. \end{cases}$$

By using the above equation and by simple induction arguments, it is possible to show that System-Z can be written in canonical form as

$$z_k(i+1) = (1 + \delta)^{k-1} \left(\delta \sum_{j=1}^r z_j(i) + \hat{\lambda} \left(\frac{r \tilde{\beta}}{4v\varepsilon} + a \right) \right) - \sum_{j=1}^{k-1} \delta (1 + \delta)^{k-j-1} z_j(i), \quad k \in \{1, \dots, r\}.$$

Hence, System-Z can be written in compact form as

$$z(i+1) = \left(A + B(\varepsilon) \right) z(i) + c(\varepsilon),$$

where A is a matrix in $\mathbb{R}^{r \times r}$ and is independent of ε , $B(\varepsilon) \in \mathbb{R}^{r \times r}$ and each element of $B(\varepsilon)$, say $b(\varepsilon)_{kj}$, has the property $\lim_{\varepsilon \rightarrow 0^+} b(\varepsilon)_{kj} = 0$ for all $k, j \in \{1, \dots, r\}$, and $c(\varepsilon) \in \mathbb{R}^{r \times 1}$ (notice that $c(\varepsilon)$ is well defined for every $\varepsilon > 0$). It is easy to see that the entries of matrix A are

$$a_{kj} = \begin{cases} \hat{\rho} \left[(1 + \hat{\rho})^{k-1} - (1 + \hat{\rho})^{k-j-1} \right], & \text{for } 1 \leq j \leq k-1, \\ \hat{\rho} (1 + \hat{\rho})^{k-1}, & \text{otherwise,} \end{cases}$$

where $\hat{\rho} \doteq \hat{\lambda}\bar{s}$. Notice that, since $\hat{\rho} > 0$, we have $|a_{kj}| = a_{kj}$ for all k and j . Then, for each $k \in \{1, \dots, r\}$, we have

$$\sum_{j=1}^r |a_{kj}| = \sum_{j=1}^r a_{kj} = r\hat{\rho}(1 + \hat{\rho})^{k-1} - \hat{\rho} \sum_{j=1}^{k-1} (1 + \hat{\rho})^{k-j-1} = (1 + \hat{\rho})^{k-1} \underbrace{(r\hat{\rho} - 1)}_{<0} + 1 < 1,$$

since $\hat{\rho} = \lambda\bar{s}/r$ and $\varrho = \lambda\bar{s} < 1$ by assumption. Therefore, the L^∞ -induced norm of matrix A satisfies: $\|A\|_\infty = \max_k \sum_{j=1}^r |a_{kj}| < 1$. Since any induced norm $\|\cdot\|$ satisfies the inequality $\rho(A) \leq \|A\|$, where $\rho(A)$ is the spectral radius of A , we conclude that $A \in \mathbb{R}^{r \times r}$ has eigenvalues strictly inside the unit disk (i.e., A is a stable matrix). Since the eigenvalues of a matrix depend continuously on the matrix entries, there exists a sufficiently small $\varepsilon > 0$ such that the matrix $A + B(\varepsilon)$ has eigenvalues strictly inside the unit disk. Accordingly, having selected a sufficiently small ε , each solution $i \mapsto z(i) \in \mathbb{R}_{\geq 0}^r$ of System-Z converges exponentially fast to the unique equilibrium point

$$z^*(\varepsilon) = \left(I_r - A - B(\varepsilon) \right)^{-1} c(\varepsilon), \quad (15)$$

where I_r is the identity matrix of size r . Combining equation (14) with the previous statement, we see that the solutions $i \mapsto \bar{n}(i)$ (where $\bar{n}(i) \doteq (\bar{n}_i^1, \dots, \bar{n}_i^r)$) and $i \mapsto y(i)$ are bounded. Thus

$$\limsup_{i \rightarrow +\infty} \bar{n}(i) \leq \limsup_{i \rightarrow +\infty} y(i) < +\infty. \quad (16)$$

Finally, we turn our attention to the third issue, namely the computation of $y \doteq \limsup_{i \rightarrow +\infty} y(i)$. Taking the lim sup of the left- and right-hand sides of equation (11), and noting that

$$\limsup_{i \rightarrow +\infty} \sqrt{y_j(i)} = \sqrt{\limsup_{i \rightarrow +\infty} y_j(i)}, \quad \text{for } j \in \{1, \dots, r\},$$

since $\sqrt{\cdot}$ is continuous and strictly monotone increasing on $\mathbb{R}_{>0}$, we obtain that

$$y_k = \hat{\lambda}a + \hat{\lambda} \sum_{j=1}^r \left(\bar{s}y_j + \frac{\tilde{\beta}}{v} \sqrt{y_j} \right); \quad (17)$$

therefore we have $y_k = y_j$ for all $k, j \in \{1, \dots, r\}$. Substituting in equation (17) and solving for y_k we obtain

$$y_k = \frac{1}{4} \left(\frac{\lambda\tilde{\beta}}{v(1-\varrho)} + \sqrt{\frac{\lambda^2\tilde{\beta}^2}{v^2(1-\varrho)^2} + \frac{4\lambda a}{1-\varrho}} \right)^2.$$

Recalling that this analysis holds for every $\varrho \in [\bar{\varrho}, 1)$, by taking ϱ arbitrarily close to 1 we obtain

$$y_k = \frac{\lambda^2\tilde{\beta}^2}{v^2(1-\varrho)^2} \quad \text{as } \varrho \rightarrow 1^-. \quad (18)$$

(Formally, we should have written $\lim_{\varrho \rightarrow 1^-} y_k(1-\varrho)^2 = \lambda^2\tilde{\beta}^2/(v^2(1-\varrho)^2)$, however, as discussed in remark 3.2, we prefer to adhere to the less formal but more transparent style of equation (18)).

Noting that from equation (16), $\limsup_{i \rightarrow +\infty} \bar{n}_i^k \leq y_k$, and recalling that $\tilde{\beta} = \beta_{\text{TSP},2} \int_{\mathcal{Q}} f^{1/2}(x) dx / \sqrt{r} + \zeta$ where ζ is an arbitrarily small constant, we obtain the desired result. ■

We are now in a position to prove Theorem 4.2

Proof of Theorem 4.2: Define $\bar{C}^k \doteq \limsup_{i \rightarrow \infty} \mathbb{E}[C_i^k]$; then we have, by using the upper bound on $\mathbb{E}[C_i^k]$ in equation (10) and neglecting constant terms,

$$\bar{C}^k \doteq \limsup_{i \rightarrow \infty} \mathbb{E}[C_i^k] \leq \frac{\tilde{\beta}}{v} \sum_{j=1}^r \sqrt{\bar{n}^j} + \bar{s} \sum_{j=1}^r \bar{n}^j \leq \frac{r\lambda\tilde{\beta}\beta}{v^2(1-\varrho)} + \frac{r\varrho\lambda\beta^2}{v^2(1-\varrho)^2}.$$

Hence, in the limit $\varrho \rightarrow 1^-$, we have

$$\bar{C}^k \leq \frac{r \lambda \beta^2}{v^2 (1 - \varrho)^2} = \frac{\lambda \beta_{\text{TSP},2}^2 \left[\int_Q f^{1/2}(x) dx \right]^2}{v^2 (1 - \varrho)^2}.$$

Consider a random demand that arrives in subregion k . Its expected steady-state system time, \bar{T}^k , will be upper bounded, as $\varrho \rightarrow 1^-$, by

$$\bar{T}^k \leq \frac{1}{2} \bar{C}^k + \frac{1}{2} \bar{s} \bar{n}^k \leq \frac{1}{2} \bar{C}^k + \frac{1}{2} \frac{\lambda \beta^2}{v^2 (1 - \varrho)^2},$$

where we used the fact that, as $\varrho \rightarrow 1^-$, the travel time along a TSP tour is negligible compared to the on-site service time requirement. Unconditioning, we obtain the claim

$$\bar{T} \leq \left(1 + \frac{1}{r}\right) \frac{\lambda \beta_{\text{TSP},2}^2 \left[\int_Q f^{1/2}(x) dx \right]^2}{2v^2 (1 - \varrho)^2}.$$

■

C. Discussion

The DC policy is optimal in light load (Theorem 4.1); moreover, if we let $r \rightarrow \infty$, the DC policy achieves the heavy load lower bound (4) for unbiased policies (Theorem 4.2). Therefore the DC policy is *both* optimal in light load *and* arbitrarily close to optimality in heavy load, and stabilizes the system in every load condition (since it stabilizes the system in the limit $\varrho \rightarrow 1^-$). Notice that with $r = 10$ the DC policy is already guaranteed to be within 10% of the optimal (for unbiased policies) performance in heavy load.

If one chooses $r > 1$, the computation of a simultaneously equitable r -partition of \mathcal{Q} requires the knowledge of f . Hence, if $r \rightarrow +\infty$, the policy is (i) provably optimal both in light and heavy load, and (ii) adaptive with respect to arrival rate λ , expected on-site service \bar{s} , and vehicle's velocity v . If, instead, $r = 1$, the policy is (i) provably optimal in light load and within a factor 2 of the optimal in heavy load, and (ii) adaptive with respect to arrival rate λ , expected on-site service \bar{s} , vehicle's velocity v , and spatial density f ; in other words, when $r = 1$, the DC policy adapts to *all* problem data. (These results should be compared with the discussion in Section III about the non-adaptive nature of the UTSP policy, which, for example, might become unstable if λ increases.)

An approximate algorithm to compute, with *arbitrarily small* error, an r -partition that is simultaneously equitable with respect to two measures can be found in [13] (specifically, see discussion on page 621); in the particular case when the density f is uniform, the problem of finding an r -partition that is simultaneously equitable with respect to f and $f^{1/2}$ becomes trivial. Given a simultaneously equitable r -partition of \mathcal{Q} , the subregions can be indexed using the following procedure: (i) a TSP tour through the medians of the subregions is computed, (ii) subregions are indexed according to the order induced by this TSP tour. In practice, the DC policy would be implemented by allowing the vehicle to skip subregions where no demand is present and by directly moving from one TSP tour to the next one, i.e., by skipping the median of subregion \mathcal{Q}_k in line 6 of Algorithm 1 (the ‘‘shortcuts’’ were not included in the presentation of the DC policy to make its analysis easier).

Note that if we assume that f is known (as it must be the case for the implementation of the DC policy with $r > 1$), then instead of using the empirical median \tilde{P}_1^* one should directly use the exact median P_1^* (which is the solution of a strictly convex problem).

The DC policy requires on-line solutions of possibly large TSPs and, therefore, practical implementations of this policy should rely on heuristics for TSPs, such as Lin-Kernighan's or Christofides' (see Section II-D). We will further discuss computation times in Section VIII.

Finally, the DC policy with $r = 1$ is similar to the generation policy presented in [27]; in particular, the generation policy has the same performance guarantees of the DC policy with $r = 1$. However, the generation policy is analyzed only for the case of *uniform* spatial density f , and its implementation *does* require the knowledge of f (while, as discussed before, the DC policy with $r = 1$ adapts to all problem data, including f). The part of Algorithm 1 in lines 5-9 is similar to the PART-TSP policy presented in [26]; however, the PART-TSP policy is only *informally* analyzed by assuming steady-state conditions, in particular no proof of stability and convergence is provided.

In the next section we present and analyze another single-vehicle policy, the Receding Horizon (RH) policy, that applies ideas of receding-horizon control to dynamic vehicle routing problems.

V. THE SINGLE-VEHICLE RECEDING HORIZON POLICY

Next, we describe the Receding Horizon (RH) policy in Algorithm 2, where D is the set of outstanding demands waiting for service and p is the current position of the vehicle; moreover, given a tour T of D , a fragment of T is a connected subgraph of T . If $D \neq \emptyset$, the vehicle selects a random η -fragment of the TSP tour through the

Algorithm 2: Receding Horizon (RH) Policy

input: Scalar $\eta \in (0, 1)$

- 1 **if** $D = \emptyset$ **then**
- 2 Let \tilde{P}_1^* be the point minimizing the sum of distances to demands serviced in the past (the initial condition for \tilde{P}_1^* is a random point in \mathcal{Q}); if $p \neq \tilde{P}_1^*$, move to \tilde{P}_1^* , otherwise stop.
- 3 **else**
- 4 **repeat**
- 5 Compute the TSP tour through all demands in D .
- 6 Uniformly randomly choose, independently from the past, a fragment of length η TSP(D) of such TSP tour.
- 7 Service the selected fragment starting from the endpoint closest to the current position.
- 8 **until** $D = \emptyset$
- 9 Repeat.

demands in D (i.e., a fragment of length η TSP(D) of such tour); note that the horizon is not fixed, but it is adjusted according to the cost of servicing the outstanding demands. In general, the performance of the system will

depend on the choice of the parameter η , which will be discussed after the analysis of the policy. Notice that the RH policy is an *unbiased* policy.

A. Stability and Performance of the RH Policy

The RH policy is attractive since it can be implemented without any knowledge of the underlying demand generation process (in particular, without any knowledge of f). On the other hand, an important caveat of receding horizon control is that closed-loop stability is not guaranteed in general. In this section, we study the stability of the RH policy and we discuss its performance.

1) *Stability of the RH Policy:* In general, the stability properties of receding horizon controllers deteriorate as the horizon becomes shorter; remarkably, the RH policy is stable in every load for every $0 < \eta < 1$.

Theorem 5.1: As $\varrho \rightarrow 1^-$, the system time for the RH policy satisfies

$$\bar{T}_{\text{RH}}(\eta) \leq \frac{\lambda \beta_{\mathcal{Q},2}^2 |\mathcal{Q}|}{v^2 (1 - \varrho)^2}, \quad \text{for all } \eta \in (0, 1),$$

where $|\mathcal{Q}|$ is the area of \mathcal{Q} and $\beta_{\mathcal{Q},2}$ is the constant appearing in equation (2).

Proof: By following the same arguments as in Lemma 4.3, it is easy to show that under the RH policy, after a transient, the number of demands at the instants when a new TSP tour is computed is very large with high probability. Then, in heavy load, we have $D \neq \emptyset$ with high probability, and the RH policy *reduces* to lines 4-8 with the condition $D = \emptyset$ always false.

We refer to the time instant t_i , $i \geq 0$, in which the vehicle computes a new TSP tour 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 and we refer to its length as the cost C_i . Let n_i be the number of outstanding demands at epoch i ; with a slight abuse of notation, we denote by $\text{TSP}(n_i)$ the length of the TSP tour through the outstanding demands at epoch i .

Since an η -fragment is randomly chosen, the expected number of demands left unserved during iteration i is $(1 - \eta) \mathbb{E}[n_i]$; then,

$$\mathbb{E}[n_{i+1}] = \underbrace{\lambda \mathbb{E}[C_i]}_{\text{newly arrived demands}} + \underbrace{(1 - \eta) \mathbb{E}[n_i]}_{\text{demands left unserved during iteration } i}.$$

The expected number of demands that receive service during iteration i is $\eta \mathbb{E}[n_i]$. Moreover, the time to reach the starting endpoint of the selected fragment is bounded by $\text{diam}(\mathcal{Q})/v$. Then, the expected time-length of iteration i can be upper bounded as

$$\mathbb{E}[C_i] \leq \text{diam}(\mathcal{Q})/v + \frac{\eta}{v} \mathbb{E}[\text{TSP}(n_i)] + \bar{s} \eta \mathbb{E}[n_i] \leq \text{diam}(\mathcal{Q})/v + \frac{\eta}{v} \beta_{\mathcal{Q},2} \sqrt{|\mathcal{Q}|} \sqrt{\mathbb{E}[n_i]} + \bar{s} \eta \mathbb{E}[n_i], \quad (19)$$

where we have applied the deterministic bound in equation (2) and Jensen's inequality for concave functions. Therefore, we obtain

$$\mathbb{E}[n_{i+1}] \leq \lambda \text{diam}(\mathcal{Q})/v + \lambda \frac{\eta}{v} \beta_{\mathcal{Q},2} \sqrt{|\mathcal{Q}|} \sqrt{\mathbb{E}[n_i]} + \varrho \eta \mathbb{E}[n_i] + (1 - \eta) \mathbb{E}[n_i].$$

By using the same techniques as in the proof of Lemma 4.4 (in particular, the eigenvalue of the ‘‘bounding’’ discrete-time linear system is $\mu = 1 - \eta(1 - \varrho)$, whose magnitude is strictly less than one for all $\eta \in (0, 1)$), it is

possible to show that

$$\bar{n} \doteq \limsup_{i \rightarrow \infty} \mathbb{E}[n_i] \leq \frac{1}{4} \left(\frac{\lambda \beta_{\mathcal{Q},2} |\mathcal{Q}|}{v(1-\varrho)} + \sqrt{\frac{\lambda^2 \beta_{\mathcal{Q},2}^2 |\mathcal{Q}|}{v^2 (1-\varrho)^2} + \frac{4\lambda \text{diam}(\mathcal{Q})}{v\eta(1-\varrho)}} \right)^2, \quad \text{for all } \eta \in (0, 1).$$

From Little's law (see, e.g., [25]), $\bar{T}_{\text{RH}} = \bar{n}/\lambda + \bar{s}$, and thus for all $\eta \in (0, 1)$

$$\bar{T}_{\text{RH}} \leq \frac{\lambda \beta_{\mathcal{Q},2}^2 |\mathcal{Q}|}{v^2 (1-\varrho)^2} \quad \text{as } \varrho \rightarrow 1^-,$$

where we neglected the terms $4\lambda \text{diam}(\mathcal{Q})/(v\eta(1-\varrho))$ and \bar{s} because of the heavy load assumption. \blacksquare

Since, by Theorem 5.1, the RH policy stabilizes the system in the limit $\varrho \rightarrow 1^-$, it also stabilizes the system for every load factor ϱ .

2) *Performance of the RH Policy:* In light load, the RH policy becomes identical to the DC policy, and therefore it is *optimal* (see Theorem 4.1). For the case $\varrho \rightarrow 1^-$, because of the dependencies that arise among demands' locations, we were unable to obtain upper bounds on the system time that are both tight and rigorous. However, we now present an heuristic analysis of the RH policy that provides interesting insights on its behavior and suggests bounds on its heavy load system time that are well matched by simulation results.

The service of an η -fragment introduces *dependencies* among the locations of the demands that are left unserved; therefore, even though the number of demands, as $\varrho \rightarrow 1^-$, becomes large with probability 1, the result in equation (1) (that requires a set of *independently* distributed demands) formally does *not* hold. However, due to the randomized selection of the η -fragments, it is quite natural to conjecture that equation (1) still provides a good estimate on the average lengths of the TSP tours that are computed. The rationale behind this conjecture is the following: When η is close to 1, most of the outstanding demands are serviced during one iteration, and the next TSP tour is mostly through newly arrived demands, that are independent by the assumptions on the demand generation process; when, instead, η is close to zero, the randomized selection of *short* fragments only introduce ‘‘negligible’’ correlations.

The above observations motivate us to use $\beta_{\text{TSP},2} \int_{\mathcal{Q}} f^{1/2}(x) dx$ (as it would be dictated by equation (1)) in equation (19), instead of $\beta_{\mathcal{Q},2} \sqrt{|\mathcal{Q}|}$ (as it is dictated by equation (2)). Then, by repeating the same steps as in the proof of Theorem 4.2, Lemma 4.4, and Theorem 5.1, we would obtain the following result for all $\eta \in (0, 1)$

$$\bar{T}_{\text{RH}}(\eta) \leq \frac{\lambda \beta_{\text{TSP},2}^2 \left[\int_{\mathcal{Q}} f^{1/2}(x) dx \right]^2}{v^2 (1-\varrho)^2}, \quad \text{as } \varrho \rightarrow 1^-. \quad (20)$$

The upper bound in equation (20) is generally tighter than the upper bound in Theorem 5.1 (notice that by Jensen's inequality for concave functions $\int_{\mathcal{Q}} f^{1/2}(x) dx \leq \sqrt{|\mathcal{Q}|} \left[\int_{\mathcal{Q}} f(x) dx \right]^{1/2} = \sqrt{|\mathcal{Q}|}$), but unlike the upper bound in Theorem 5.1, it has not been established formally. Simulation results (see Section VIII) indeed confirm the upper bound in equation (20).

Finally, note that when η is close to zero, the RH policy is conceptually similar to the DC policy with $r \rightarrow \infty$, so we also speculate that in heavy load the performance of the RH policy improves as the horizon becomes shorter. Simulation results (see Section VIII) indeed confirm this behavior.

B. Discussion

The RH policy is optimal in light load and stabilizes the system in every load condition (these two statements have been established rigorously). The implementation of the RH policy does *not* require the knowledge of λ, \bar{s}, v and f . Therefore, the RH policy *adapts* to arrival rate λ , expected on-site service \bar{s} , vehicle's velocity v , and spatial density f ; in other words, the RH policy adapts to *all* problem data.

In Section IV we saw that also the DC policy with $r = 1$ adapts to all problem data. Simulation results (see Section VIII) show that in heavy load the RH policy with a short horizon η (say, $\eta \approx 0.2$) performs better than the DC policy with $r = 1$ (the light load behavior is clearly the same). Therefore, to date, the best available policy for the 1-DTRP that is unbiased and adapts to all problem data seems to be the RH policy with $\eta \approx 0.2$.

It is natural to wonder how the RH policy performs if the η -fragment is not selected randomly, but it is instead selected according to the rule: Find the fragment of length $\eta \text{TSP}(D)$ that maximizes the number of reached demands. In other words, the vehicle now looks for a maximum-reward η -fragment. Clearly, Theorem 5.1 still applies. Simulation results (see [32]) show that this variant of the RH policy appears to perform better than an optimal unbiased policy but not as well as an optimal biased policy.

As for the DC policy, the RH policy requires on-line solutions of possibly large TSPs and, therefore, practical implementations of this policy should also rely on heuristics for TSPs, such as Lin-Kernighan's or Christofides'. We will further discuss computation times in Section VIII.

VI. ADAPTIVE AND DISTRIBUTED POLICIES FOR THE m -DTRP

In this section we extend the previous single-vehicle policies to the multi-vehicle case. First, we prove that in heavy load, given a single-vehicle, unbiased optimal policy π^* and m vehicles, there exists an unbiased π^* -partitioning policy that is optimal. This result and the optimality of the SQM policy (which is a partitioning policy that is optimal in light load) lead us to propose, for the multi-vehicle case, *partitioning policies* that use the DC policy or the RH policy as single-vehicle policies. Finally, we discuss how a partition can be computed in a distributed fashion, so that the proposed policies are amenable to distributed implementation.

A. Optimality of Partitioning Policies in Heavy Load

The following theorem holds.

Theorem 6.1: As $\varrho \rightarrow 1^-$, given a single-vehicle, unbiased optimal policy π^* and m vehicles, a π^* -partitioning policy using an m -partition which is simultaneously equitable with respect to f and $f^{1/2}$ is an optimal unbiased policy for the m -DTRP.

Proof: Let π^* be an optimal unbiased policy (e.g., the DC policy with $r \rightarrow \infty$). Construct an m -partition $\{\mathcal{Q}_k\}_{k=1}^m$ of \mathcal{Q} that is simultaneously equitable with respect to f and $f^{1/2}$ (such partition is guaranteed to exist as discussed in Section II); assign one vehicle to each subregion. Each vehicle executes the single-vehicle unbiased policy π^* to service demands that fall within its own subregion. The probability that a demand falls in subregion \mathcal{Q}_k is equal to $\int_{\mathcal{Q}_k} f(x) dx = 1/m$. Notice that the arrival rate λ_k in subregion \mathcal{Q}_k is reduced by a factor $\int_{\mathcal{Q}_k} f(x) dx$,

i.e., $\lambda_k = \lambda \int_{\mathcal{Q}_k} f(x) dx = \lambda/m$; therefore, the load factor in subregion \mathcal{Q}_k (where only one vehicle provides service) is $\varrho_k = \lambda_k \bar{s} = \lambda \bar{s}/m = \varrho < 1$, in particular the necessary condition for stability *is satisfied* in every subregion. Finally, given that a demand falls in subregion \mathcal{Q}_k , the conditional density for its location (whose support is \mathcal{Q}_k) is $f(x)/\left(\int_{\mathcal{Q}_k} f(x) dx\right) = m f(x)$. Then, by the law of total expectation we have

$$\begin{aligned} \bar{T} &= \sum_{k=1}^m \left(\int_{\mathcal{Q}_k} f(x) dx \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda_k}{v^2 (1 - \varrho_k)^2} \left[\int_{\mathcal{Q}_k} \sqrt{\frac{f(x)}{\int_{\mathcal{Q}_k} f(x) dx}} dx \right]^2 \right) \\ &= \sum_{k=1}^m \left(\frac{1}{m} \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda}{v^2 (1 - \varrho)^2} \left[\int_{\mathcal{Q}_k} f^{1/2}(x) dx \right]^2 \right) \\ &= \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda \left[\int_{\mathcal{Q}} f^{1/2}(x) dx \right]^2}{m^2 v^2 (1 - \varrho)^2}, \quad \text{as } \varrho \rightarrow 1^-. \end{aligned} \quad (21)$$

Thus, the bound in (4) is achieved. We finally show that such multi-vehicle policy is indeed unbiased. In fact, assume steady-state conditions, and let X be the location of a randomly chosen demand, T be its system time and \mathcal{S} be an arbitrary subset of \mathcal{Q} ; then, by the law of total expectation we have

$$\mathbb{E}[T | X \in \mathcal{S}] = \mathbb{E}\left[\mathbb{E}[T | X \in \mathcal{S} \cap \mathcal{Q}_k] | X \in \mathcal{S}\right].$$

Since π^* is unbiased, it must hold $\mathbb{E}[T | X \in \mathcal{S} \cap \mathcal{Q}_k] = \mathbb{E}[T | X \in \mathcal{Q}_k]$; moreover, from the previous analysis, we have $\mathbb{E}[T | X \in \mathcal{Q}_k] = \bar{T}_{\pi^*}$, which is a constant independent of \mathcal{Q}_k . Hence, we finally obtain $\mathbb{E}[T | X \in \mathcal{S}] = \bar{T}_{\pi^*}$. The claim then follows from the fact that \mathcal{S} was chosen arbitrarily. ■

Remark 6.2: Some remarks are in order.

- 1) Theorem 6.1 proves Conjecture 1 (made only for uniform spatial densities f) in [11].
- 2) A similar result can be proven for biased policies; in this case the m -partition should be simultaneously equitable with respect to f and $f^{2/3}$.
- 3) The proof of Theorem 6.1 relies on m -partitions that are simultaneously equitable with respect to f and $f^{1/2}$. We next show that a π -partitioning policy using an m -partition which is equitable with respect to $f^{1/2}$ may be unstable; moreover, we show that, in general, a π^* -partitioning policy using an m -partition which is equitable with respect to f does not achieve the optimal unbiased performance, however it is always within a factor m of it.

Theorem 6.3: Given a single-vehicle policy π and m vehicles, a π -partitioning policy using an m -partition which is equitable with respect to $f^{1/2}$ may be unstable for values of ϱ strictly less than one.

Proof: The arrival rate in subregion \mathcal{Q}_k , $k \in \{1, \dots, m\}$, is $\lambda_k = \lambda \int_{\mathcal{Q}_k} f(x) dx$, and therefore the load factor in subregion \mathcal{Q}_k (where only one vehicle provides service) is $\varrho_k = \lambda \int_{\mathcal{Q}_k} f(x) dx \bar{s}$. In general, an m -partition that is equitable with respect to $f^{1/2}$ is *not* equitable with respect to f . If this is the case, there exists $\bar{k} \in \{1, \dots, m\}$ and $\varepsilon > 0$ such that $\int_{\mathcal{Q}_{\bar{k}}} f(x) dx = 1/m + \varepsilon$; then, $\varrho_{\bar{k}} = \varrho + \varepsilon \lambda \bar{s}$. Hence, for every value of ϱ such that $1 - \varepsilon \lambda \bar{s} \leq \varrho < 1$, such multi-vehicle policy is unstable (since in subregion $\mathcal{Q}_{\bar{k}}$ we have $\varrho_{\bar{k}} \geq 1$), even though the load factor ϱ is strictly less than one.

Theorem 6.4: Let π be an unbiased policy such that in heavy load $\bar{T}_\pi/\bar{T}_U^* \leq \gamma$. As $\rho \rightarrow 1^-$, given the policy π and m vehicles, a π -partitioning policy using an m -partition which is equitable with respect to f does not achieve, in general, the optimal unbiased performance, however it is always within a factor γm of it.

Proof: Let π be an unbiased policy. The environment \mathcal{Q} is partitioned into m subregions \mathcal{Q}_k , $k \in \{1, \dots, m\}$, such that $\int_{\mathcal{Q}_k} f(x) dx = 1/m$ for all k , and one vehicle is assigned to each subregion. Each vehicle executes the single-vehicle policy π to service demands that fall within its own subregion. Similarly to the proof of Theorem 6.1 we have, by the law of total expectation,

$$\begin{aligned} \bar{T} &\geq \sum_{k=1}^m \left(\int_{\mathcal{Q}_k} f(x) dx \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda_k}{v^2(1-\rho_k)^2} \left[\int_{\mathcal{Q}_k} \sqrt{\frac{f(x)}{\int_{\mathcal{Q}_k} f(x) dx}} dx \right]^2 \right) \\ &= \sum_{k=1}^m \left(\frac{1}{m} \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda}{v^2(1-\rho)^2} \left[\int_{\mathcal{Q}_k} f^{1/2}(x) dx \right]^2 \right) \quad \text{as } \rho \rightarrow 1^-. \end{aligned} \quad (22)$$

Define the vector

$$z = \left(\int_{\mathcal{Q}_1} f^{1/2}(x) dx, \int_{\mathcal{Q}_2} f^{1/2}(x) dx, \dots, \int_{\mathcal{Q}_m} f^{1/2}(x) dx \right) \in \mathbb{R}^m;$$

we will denote by $\|z\|_2$ and $\|z\|_1$ the 2-norm and the 1-norm, respectively, of vector z . Notice that $\|z\|_1 = \int_{\mathcal{Q}} f^{1/2}(x) dx$. Since for any vector $z \in \mathbb{R}^m$ it holds $\|z\|_1 \leq \sqrt{m}\|z\|_2$, equation (22) becomes

$$\begin{aligned} \bar{T} &\geq \frac{1}{m} \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda}{v^2(1-\rho)^2} \|z\|_2^2 \geq \frac{1}{m} \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda}{v^2(1-\rho)^2} \frac{\|z\|_1^2}{m} \\ &= \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda \left[\int_{\mathcal{Q}} f^{1/2}(x) dx \right]^2}{m^2 v^2 (1-\rho)^2} = \bar{T}_U^*, \quad \text{as } \rho \rightarrow 1^-, \end{aligned}$$

where \bar{T}_U^* is the optimal unbiased performance. In general, an m -partition that is equitable with respect to f is *not* equitable with respect to $f^{1/2}$; therefore, in general, the inequality $\|z\|_1 \leq \sqrt{m}\|z\|_2$ is strict and \bar{T} is larger than the optimal unbiased performance.

On the other hand, for any vector $z \in \mathbb{R}^m$ it also holds $\|z\|_2 \leq \|z\|_1$, therefore by using the same reasoning as in equation (22) we obtain

$$\begin{aligned} \bar{T} &\leq \frac{1}{m} \gamma \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda}{v^2(1-\rho)^2} \|z\|_2^2 \leq \frac{1}{m} \gamma \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda}{v^2(1-\rho)^2} \|z\|_1^2 \\ &= \gamma \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda \left[\int_{\mathcal{Q}} f^{1/2}(x) dx \right]^2}{m v^2 (1-\rho)^2} = \gamma m T_U^*, \quad \text{as } \rho \rightarrow 1^-, \end{aligned}$$

where, again, \bar{T}_U^* is the optimal unbiased performance. ■

Note that in this case the resulting multi-vehicle policy might not be unbiased.

Remark 6.5: A similar result holds for biased policies: in this case the 3-norm is the relevant norm to use.

B. Distributed Policies for the m -DTRP and Discussion

The optimality of the SQM policy and Theorem 6.1 suggest the following distributed multi-vehicle version of the DC policy:

- 1) the vehicles compute in a distributed way an m -median of \mathcal{Q} that induces a Voronoi partition that is equitable with respect to f and $f^{1/2}$,
- 2) the vehicles assign themselves to the subregions (thus, there is a one-to-one correspondence between vehicles and subregions),
- 3) each vehicle executes the single-vehicle DC policy to service demands that fall within its own subregion, by using the median of the subregion instead of \tilde{P}_1^* .

(In an identical way, we can obtain a multi-vehicle version of the RH policy.)

If there exists, given \mathcal{Q} and f , an m -median of \mathcal{Q} that induces a Voronoi partition that is equitable with respect to f and $f^{1/2}$, then the above policy is optimal both in light load and arbitrarily close to optimality in heavy load, and stabilizes the system in every load condition. There are two main issues with the above policy, namely (i) existence of an m -median of \mathcal{Q} that induces a Voronoi partition that is equitable with respect to f and $f^{1/2}$, and (ii) how to compute it in a distributed way. In [33], we showed that for some choices of \mathcal{Q} and f a median Voronoi diagram that is equitable with respect to f and $f^{1/2}$ fails to exist. On the other hand, in [34], [35], we presented a synchronous, spatially-distributed partitioning algorithm that, for any possible choice of \mathcal{Q} and f , provides a partition that is equitable with respect to f and represents a “good” approximation of a median Voronoi diagram (see [35] for details on the metrics that we use to judge “closeness” to median Voronoi diagrams; moreover, in [35], the motion toward the “centroid” should be replaced with a motion toward the median). Furthermore, if an m -median of \mathcal{Q} that induces a Voronoi partition that is equitable with respect to f exists, the algorithm will locally converge to it.

Accordingly, we define the multi-vehicle Divide & Conquer policy as follows

Algorithm 3: Multi-Vehicle DC (m -DC) Policy

- 1 The vehicles run the spatially-distributed partitioning algorithm defined as a set of differential equations in [35] (replacing the motion toward the centroid with a motion toward the median, and where the “power generators” are just *artificial* variables locally controlled by each vehicle, see [35] for the details).
 - 2 The vehicles assign themselves to the subregions (this part is indeed a by-product of the algorithm in [35]).
 - 3 Simultaneously, each vehicle executes the single-vehicle DC policy inside its own subregion.
-

According to Theorem 6.4 the m -DC policy is within a factor $(1 + 1/r)m$ of the optimal unbiased performance in heavy load (since the algorithm in [35] *always* provides a partition that is equitable with respect to f), and stabilizes the system in every load condition. In general, the m -DC policy is only *suboptimal* in light load; note, however, that the computation of the global minimum of the Weber function H_m (which is non-convex for $m > 1$)

is difficult for $m > 1$ (it is NP-hard for the discrete version of the problem); therefore, for $m > 1$, suboptimality has also to be expected from *any* practical implementation of the SQM policy. If an m -median of \mathcal{Q} that induces a Voronoi partition that is equitable with respect to f exists, the partitioning algorithm used by the m -DC will locally converge to it, thus we say that the m -DC policy is “locally” optimal in light load.

Note that, when the density is uniform, a partition that is equitable with respect to f is also equitable with respect to $f^{1/2}$; therefore, when the density is uniform the m -DC policy is arbitrarily close to optimality in heavy load (see Theorem 6.1).

The m -DC policy *adapts* to arrival rate λ , expected on-site service \bar{s} , and vehicle’s velocity v ; however, it requires the knowledge of f . The key feature of the partitioning algorithm in [35] is that it is *spatially-distributed* (see, e.g., [36] for a rigorous definition of spatially-distributed algorithms); therefore, the m -DC policy is indeed a spatially-distributed control policy.

The multi-vehicle version of the RH policy (which we call m -RH policy) can be defined in an identical way. Its properties are summarized in Table II in Section I.

VII. DISTRIBUTED POLICIES FOR THE CASE WITH ZERO ON-SITE SERVICE TIME

In the formulation of the DTRP we assumed that $\bar{s} > 0$ (see Section III). However, it is of interest to also study the case $\bar{s} = 0$, since in some applications the on-site service time might actually be zero (or negligible). We show that, perhaps surprisingly, when $\bar{s} = 0$ a π^* -partitioning policy using an m -partition which is equitable with respect to $f^{1/2}$ is an optimal unbiased policy for the m -DTRP. Hence, when $\bar{s} = 0$, optimality does *not* require simultaneous equitability, a considerable simplification.

Specifically, when $\bar{s} = 0$, the light load regime is defined as $\lambda \rightarrow 0^+$, while the heavy load regime is defined as $\lambda \rightarrow +\infty$. Then, it is quite straightforward to verify that all the performance bounds in Sections III, IV, and V hold by simply substituting $\varrho = 0$, and by formally changing the limits $\varrho \rightarrow 0^+$ and $\varrho \rightarrow 1^-$ with $\lambda \rightarrow 0^+$ and $\lambda \rightarrow +\infty$. For example, equation (4) reads

$$\bar{T}_U^* \geq \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda \left(\int_{\mathcal{Q}} f^{1/2}(x) dx \right)^2}{m^2 v^2} \quad \text{as } \lambda \rightarrow +\infty,$$

and equation (7) reads

$$\bar{T}_{\text{DC}}(r) \leq \left(1 + \frac{1}{r}\right) \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda \left[\int_{\mathcal{Q}} f^{1/2}(x) dx \right]^2}{v^2} \quad \text{as } \lambda \rightarrow +\infty,$$

where r is the number of subregions. The major difference is about which type of partition one should use in heavy load (i.e., when λ is large). We have the following result.

Theorem 7.1: Assume $\bar{s} = 0$. As $\lambda \rightarrow +\infty$, given a single-vehicle, unbiased optimal policy π^* and m vehicles, a π^* -partitioning policy using an m -partition which is equitable with respect to $f^{1/2}$ is an optimal unbiased policy for the m -DTRP.

Proof: The proof is very similar to the one of theorem 6.1. Let π^* be an optimal unbiased policy (e.g., the DC policy with $r \rightarrow \infty$). Construct an m -partition $\{\mathcal{Q}_k\}_{k=1}^m$ of \mathcal{Q} that is equitable with respect to $f^{1/2}$;

thus $\int_{\mathcal{Q}_k} f^{1/2}(x) dx = \int_{\mathcal{Q}} f^{1/2}(x) dx / m$. Notice that the arrival rate λ_k in subregion \mathcal{Q}_k is reduced by a factor $\int_{\mathcal{Q}_k} f(x) dx$, i.e., $\lambda_k = \lambda \int_{\mathcal{Q}_k} f(x) dx$. Moreover, given that a demand falls in subregion \mathcal{Q}_k , the conditional density for its location (whose support is \mathcal{Q}_k) is $f(x) / \left(\int_{\mathcal{Q}_k} f(x) dx \right)$. By the law of total expectation we have

$$\begin{aligned} \bar{T} &= \sum_{k=1}^m \left(\int_{\mathcal{Q}_k} f(x) dx \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda \int_{\mathcal{Q}_k} f(x) dx}{v^2} \left[\int_{\mathcal{Q}_k} \sqrt{\frac{f(x)}{\int_{\mathcal{Q}_k} f(x) dx}} dx \right]^2 \right) \\ &= \sum_{k=1}^m \left(\int_{\mathcal{Q}_k} f(x) dx \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda}{v^2} \left[\frac{\int_{\mathcal{Q}_k} f^{1/2}(x) dx}{m} \right]^2 \right) \\ &= \frac{\beta_{\text{TSP},2}^2}{2} \frac{\lambda \left[\int_{\mathcal{Q}} f^{1/2}(x) dx \right]^2}{m^2 v^2}, \quad \text{as } \lambda \rightarrow +\infty, \end{aligned} \quad (23)$$

where in the last equality we have used the fact that $\sum_{k=1}^m \int_{\mathcal{Q}_k} f(x) dx = 1$. Thus, the bound in (4) is achieved. We finally show that such multi-vehicle policy is indeed unbiased. In fact, assume steady-state conditions, and let X be the location of a randomly chosen demand, T be its system time and \mathcal{S} be an arbitrary subset of \mathcal{Q} ; then, by the law of total expectation we have

$$\mathbb{E}[T | X \in \mathcal{S}] = \mathbb{E} \left[\mathbb{E}[T | X \in \mathcal{S} \cap \mathcal{Q}_k] | X \in \mathcal{S} \right].$$

Since π^* is unbiased, it must hold $\mathbb{E}[T | X \in \mathcal{S} \cap \mathcal{Q}_k] = \mathbb{E}[T | X \in \mathcal{Q}_k] = \beta_{\text{TSP},2}^2 \lambda \left(\int_{\mathcal{Q}_k} f^{1/2}(x) dx \right)^2 / (2m^2 v^2)$, which is a constant independent of \mathcal{Q}_k . The claim then follows from the fact that \mathcal{S} was chosen arbitrarily. ■

Theorem 7.1 is very important from a practical viewpoint, since it means that when $\bar{s} = 0$ optimality does not require simultaneous equitability. In light of theorem 7.1, when $\bar{s} = 0$ the m -DC policy should be defined in the same way of algorithm 3, with the exception that the partitioning algorithm in [35] should partition \mathcal{Q} according to $f^{1/2}$ (instead of f). Since the partitioning algorithm in [35] is then guaranteed to converge to a partition that is equitable with respect to $f^{1/2}$, we conclude that when $\bar{s} = 0$ the spatially-distributed m -DC policy (with the modification that equitability should be with respect to $f^{1/2}$) is arbitrarily close to optimality in heavy load for *any* f .

VIII. SIMULATION EXPERIMENTS

In this section we discuss, through the use of simulations, the heavy load behavior of both the DC and the RH policy, and we comment on their relative performance.

All simulations are performed on a machine with a 2.4GHz Intel Core Duo processor and 4GB of RAM. The code is written in Matlab[®]7.4 with external calls to the program `linkern`, for which we set a running time bound of one second. In all simulations we consider a circular environment $\mathcal{Q} = \{(x, y) \in \mathbb{R}^2 | x^2 + y^2 \leq 1/\pi\}$ (hence, the area of \mathcal{Q} is 1), a vehicle's velocity $v = 1$, and an on-site service time uniformly distributed in the interval $[0, 1]$ (thus $\bar{s} = 0.5$).

The spatial density $f : \mathcal{Q} \rightarrow \mathbb{R}_+$ used in the simulation experiments is

$$f(x) = \begin{cases} \frac{1-\delta}{\epsilon}, & x \in \mathcal{Q}_1, \\ \frac{\delta}{1-\epsilon}, & x \in \mathcal{Q}_2, \end{cases}$$

where $\mathcal{Q}_1 = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq \epsilon/\pi\}$, $\mathcal{Q}_2 = \mathcal{Q} \setminus \mathcal{Q}_1$, and $\epsilon \in (0, 1)$. If $1 - \delta = \epsilon$, then the density f is uniform; if, instead, $1 - \delta > \epsilon$, then the density f has a peak in subregion \mathcal{Q}_1 . In all simulation experiments, we set $\epsilon = 0.1$.

A. Heavy-Load Performance of the DC Policy

We first consider a uniform spatial density, i.e., $\delta = 1 - \epsilon = 0.9$. We consider 5 values of the load factor, namely $\varrho = 0.9, 0.93, 0.95, 0.97, 0.99$. Circles in the left-hand figure of Fig. 1 represent the ratios between the experimental system time $\bar{T}_{\text{DC}}^*(1)$ (i.e., in the DC policy $r = 1$) and \bar{T}_{U}^* (whose expression is given in equation (4)), for the different values of ϱ . We then repeat the simulation process by executing the DC policy with $r = 16$. Squares in the left-hand figure of Fig. 1 represent the ratios between the experimental system time $\bar{T}_{\text{DC}}^*(16)$ (i.e., in the DC policy $r = 16$) and \bar{T}_{U}^* , for the different values of ϱ . It can be observed that for $r = 1$ the ratios $\bar{T}_{\text{DC}}^*(1)/\bar{T}_{\text{U}}^*$ decrease as ϱ approaches one and tend to 2, in accordance with the upper bound in Theorem 4.2 (in fact, $1 + 1/r = 2$ in this case). Similarly, for $r = 16$ the ratios $\bar{T}_{\text{DC}}^*(16)/\bar{T}_{\text{U}}^*$ decrease as ϱ approaches one and tend to $(1 + 1/16) \approx 1.06$, in accordance again with the upper bound in Theorem 4.2. These results confirm the theoretical analysis in Section IV and suggest that the upper bound (7) is indeed tight for the various values of r . Note that for all values of ϱ the experimental system time is *larger* than the upper bound; this is due to one or a combination of the following reasons: First, the upper bound formally holds only in the limit $\varrho \rightarrow 1^-$. Second, we are using an approximate solution for the optimal TSP.

In Table III we show for both $r = 1$ and $r = 16$ the average number of points (denoted by $|\text{TSP}|$) through which a TSP tour is computed and the average time required for such computation. Specifically, the computation time includes (i) the time to read the points from an input file, (ii) the TSP computation time, and (iii) the time to save the tour on a file. Recall that we implemented the DC policy by using `linkern` as a solver to generate approximations to the optimal TSP tour, and we set a running time bound of one second for this solver. One can observe that the computation times are rather small even for values of ϱ very close to one; moreover, since the experimental results are very close to the theoretical bounds, we argue that `linkern` (with a running time bound of one second) computes TSP tours very close to the optimal ones (indeed, as discussed in Section II-D, `linkern` provided tours with cost in the order of 5% of the optimal tour cost for most of the benchmarks we considered). Hence, we conclude that the DC policy can be effectively implemented in real-world applications by using `linkern` as a subroutine.

The same set of simulations is performed for a non-uniform spatial density; in particular we consider $\delta = 0.6$, i.e., a peak in the small subregion \mathcal{Q}_1 . Note that in this case a simultaneously equitable partition can be trivially obtained by using radial cuts. Results are shown in the right-hand figure of Fig. 1. Circles represent the ratios between the experimental system time and \bar{T}_{U}^* (whose expression, again, is given in equation (4)) for $r = 1$, while squares represent the ratios between the experimental system time and \bar{T}_{U}^* for $r = 16$. As before, the results are in accordance with the upper bound in Theorem 4.2.

Finally, from Fig. 1 it can be observed that as $\varrho \rightarrow 1^-$ the DC policy with $r = 16$ performs better than the

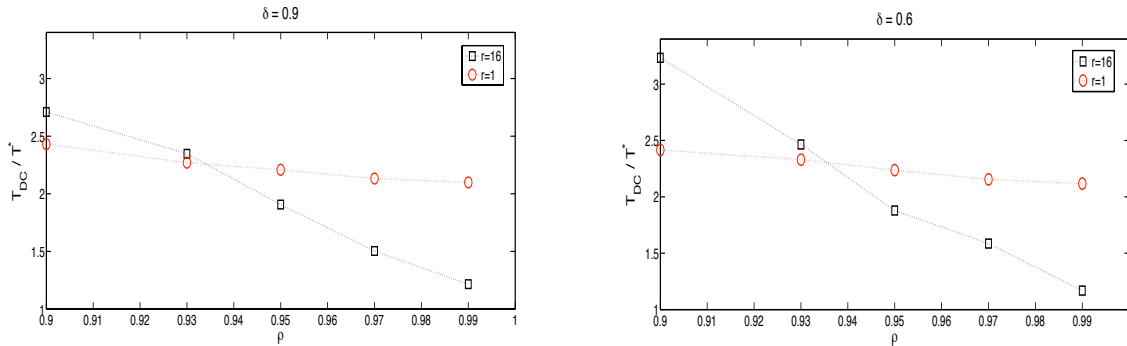


Fig. 1. Left Figure: Ratio between experimental system time under the DC policy and \bar{T}_U^* (whose expression is given in equation (4)) in the case of uniform density (i.e., $\delta = 0.9$). Right Figure: Ratio between experimental system time under the DC policy and \bar{T}_U^* in the case of non-uniform density ($\delta = 0.6$). Circles correspond to the DC policy with $r = 1$, while squares correspond to the DC policy with $r = 16$.

TABLE III
COMPUTATION TIMES FOR DC POLICY

ϱ	$r = 1$		$r = 16$	
	TSP	Time	TSP	Time
0.9	$\simeq 200$ points	$\simeq 0.3s$	$\simeq 20$ points	$\simeq 0.03s$
0.93	$\simeq 400$ points	$\simeq 0.8s$	$\simeq 60$ points	$\simeq 0.1s$
0.95	$\simeq 800$ points	$\simeq 1.1s$	$\simeq 70$ points	$\simeq 0.15s$
0.97	$\simeq 2250$ points	$\simeq 1.7s$	$\simeq 200$ points	$\simeq 0.3s$
0.99	$\simeq 21000$ points	$\simeq 5s$	$\simeq 1300$ points	$\simeq 1.3s$

DC policy with $r = 1$, precisely by a factor 2 in accordance with the upper bound in Theorem 4.2. On the other hand, for moderate values of ϱ , say $\varrho \approx 0.9$, the DC policy with $r = 1$ performs better than the DC policy with $r = 16$. This result is expected since for moderate values of ϱ and “large” values of r the number of demands in each subregion is low, and therefore equation (1) is no longer applicable (while it is still applicable for $r = 1$).

The computation times are very similar to those shown in Table III and therefore they are omitted.

B. Heavy-Load Performance of the RH Policy

We first consider a uniform spatial density. We consider 3 values of the load factor, namely $\varrho = 0.95, 0.97, 0.99$, and 7 values of η , namely $\eta = 0.1, 0.2, 0.3, 0.7, 0.8, 0.9, 1$. Formally, the RH policy is not defined for $\eta = 1$, however, by setting $\eta = 1$ in the definition of the RH policy, we simply obtain the DC policy with $r = 1$. The left-hand figure of Fig. 2 shows the ratios between the experimental $\bar{T}_{RH}(\eta)$ and \bar{T}_U^* (whose expression is given in equation (4)) for each pair (ϱ, η) . The same set of simulations is performed for a non-uniform spatial density; in particular we consider $\delta = 0.6$. Results are shown in the right-hand figure of Fig. 2. The results confirm that the system time \bar{T}_{RH} follows the $\lambda|\mathcal{Q}|/(1-\varrho)^2$ growth predicted by Theorem 5.1. In Section V we argued that in heavy load $T_{RH}(\eta) \leq 2\bar{T}_U^*$ (see equation (20)) for all $\eta \in (0, 1)$; from Fig. 2, it can be observed that as ϱ




Fig. 2. Left Figure: Ratio between experimental system time under the RH policy and \bar{T}_U^* (whose expression is given in equation (4)) in the case of uniform density (i.e., $\delta = 0.9$). Right Figure: Ratio between experimental system time under the RH policy and \bar{T}_U^* in the case of non-uniform density ($\delta = 0.6$).

approaches one this asymptotic bound seems to be confirmed, both for a uniform f and a non-uniform f . We also argued in Section V that the performance of the RH policy should improve as the horizon becomes shorter: this is indeed the case, in particular simulation results show that optimal performance is achieved for $\eta \approx 0.2$. Finally, the computation times for the RH policy with any η are very similar to those of the DC policy with $r = 1$ and therefore they are omitted.

C. Comparison between DC policy and RH policy

The RH policy should be compared with the version of the DC policy that is also adaptive with respect to f , i.e., with the DC policy with $r = 1$. Indeed, as observed before, the DC policy with $r = 1$ is the same as the RH policy when we set $\eta = 1$. Therefore, we can compare the two policies by looking at Fig. 2. One should note that the RH policy with $\eta \approx 0.2$ performs consistently better than the DC policy with $r = 1$ (which, again, is the same as the RH policy with $\eta = 1$), in particular it decreases the system time by $\approx 20\%$.

D. Performance of the Multi-Vehicle DC Policy

In this section we study, through the use of simulations, the behavior of the m -DC policy; we focus, in particular, on the heavy load scenario (in light load, the system time is directly related to the convergence of the partitioning algorithm to a median Voronoi diagram, for which we refer to [35]). When the density f is uniform, a partition that is equitable with respect to f is also equitable with respect to $f^{1/2}$. By following the same arguments as in equation (23) one can show that in heavy load the steady-state system time within each subregion is equal to the system time within the entire environment. Hence, when f is uniform and the system is in heavy load the simulation results for the m -DC policy are virtually identical to the simulation results for the 1-DC policy (presented above), and therefore they are omitted.

When the density f is *not* uniform, a partition that is equitable with respect to f is not necessarily equitable with respect to $f^{1/2}$. Theorem 6.4 however, guarantees that the m -DC policy in heavy load is within a factor $(1 + 1/r)m$

of the optimal unbiased performance. We verify this result through simulation experiments. We consider a square environment $\mathcal{Q} = \{(x, y) \in \mathbb{R}^2 \mid 0 \leq x \leq 1, 0 \leq y \leq 1\}$ (hence, the area of \mathcal{Q} is 1), vehicles' velocity $v = 1$, and an on-site service time uniformly distributed in the interval $[0, 1]$ (thus $\bar{s} = 0.5$). The spatial density used in the simulations is:

$$f(x) = \begin{cases} 3, & x \in \mathcal{Q}_1, \\ 1/3, & x \in \mathcal{Q}_2, \end{cases}$$

where $\mathcal{Q}_1 = \{(x, y) \in \mathbb{R}^2 \mid 0.5 \leq x \leq 1, 0.5 \leq y \leq 1\}$ and $\mathcal{Q}_2 = \mathcal{Q} \setminus \mathcal{Q}_1$. We consider a load factor $\rho = 0.9$ and 4 values for the number of vehicles m , namely $m = 2, 4, 6, 8$. Within each region, we simulate the DC policy with $r = 1$, which has a constant factor guarantee equal to 2 (see Theorem 4.2). The circles in figure 3 represent the ratios between the experimental system time $\bar{T}_{m\text{-DC}}$ and \bar{T}_U^* , for the different values of m . The squares in figure 3 represent the upper bounds for such ratios, which, according to Theorem 6.4, are equal to $2m$. The simulation results show that the ratios $\bar{T}_{m\text{-DC}}/\bar{T}_U^*$ seem not to increase linearly with m , or, in other words, that the upper bound in theorem 6.4 is conservative.




Fig. 3. The circles represent the ratios between the experimental system times under the m -DC policy (with $r = 1$) and \bar{T}_U^* ; the squares represent the theoretical upper bounds on these ratios. The load factor is $\rho = 0.9$.

IX. CONCLUSION

The DTRP is a widely applicable model for dynamic task-assignment problems in robotic networks. In this paper, our objective was to design unbiased policies for the DTRP that (i) are adaptive (in particular do *not* require the knowledge of the arrival rate λ and the statistics of the on-site service time), (ii) enjoy provable performance guarantees (in particular, are provably stable under any load condition), and (iii) are spatially-distributed. To this end, we introduced two new policies for the 1-DTRP, namely the DC policy and the RH policy, that were subsequently extended to the multi-vehicle case through the idea of partitioning policies. Our policies require on-line solutions of possibly large TSPs; however, we have argued in Section VIII that our policies can be effectively implemented in real-world applications by using `linkern` as a solver to generate approximations to the optimal TSP tour.

Tables I and II in the introduction provide a synoptic view of the results we achieved. We mention that our work also shows the importance of simultaneously equitable partitions in the context of robotic networks.

The focus of this paper was on the DTRP. Note, however, that in many dynamic vehicle routing problems (e.g., problems with stochastic time constraints [29], [37] and with priorities [30]) partitioning policies provide “good”

performance. Hence, following the discussion in Section VI, we argue that combining the spatially distributed partitioning algorithms presented in [35] with appropriate partitioning and single-vehicle routing policies yields adaptive and distributed algorithms for a wide variety of dynamic vehicle routing problems. The approach presented in this paper is therefore rather general and not specifically tailored to the DTRP.

This paper leaves numerous important extensions open for further research. First, we plan to devise spatially-distributed algorithms that provide simultaneously equitable partitions; such algorithms would make the m -DC policy optimal for any density f . Second, finding an unbiased policy for the 1-DTRP that is optimal in heavy load and does not rely on the knowledge of f is both practically important and theoretically interesting. Third, it is also of interest to design optimal or near-optimal policies that are *biased* in heavy load and that are adaptive and distributed (recall that allowing biased service results in a strict reduction of the optimal mean system time for any non-uniform density f). Finally, in our analysis, we considered omni-directional vehicles with first order dynamics: more realistic dynamical models will have to be taken into account for practical application to UAVs or other systems.

ACKNOWLEDGMENTS

The authors thank Stephen L. Smith for insightful feedback during the development of this work. The research leading to this work was partially supported by the National Science Foundation through grants 0705451 and 0705453, and by the Air Force Office of Scientific Research through grant AFOSR MURI-FA9550-07-1-0528.

REFERENCES

- [1] Gurdal Arslan, Jason R. Marden, and Jeff S. Shamma. Autonomous vehicle-target assignment: A game-theoretical formulation. *Journal of Dynamic Systems, Measurement, and Control*, 129(5):584–596, 2007.
- [2] B. J. Moore and K. M. Passino. Distributed task assignment for mobile agents. *IEEE Trans. on Automatic Control*, 52(4):749–753, 2007.
- [3] L. F. Bertuccelli, M. Alighanbari, and J. P. How. Robust planning for coupled cooperative UAV missions. In *Proc. IEEE Conf. on Decision and Control*, volume 3, pages 2917–2922, Bahamas, December 2004.
- [4] S. L. Smith and F. Bullo. Monotonic target assignment for robotic networks. *IEEE Trans. Automatic Control*, 54(9):2042–2057, 2009.
- [5] M. Alighanbari and J. P. How. Cooperative task assignment of unmanned aerial vehicles in adversarial environments. In *Proc. American Control Conf.*, volume 7, pages 4661–4666, Portland, Oregon, June 2005.
- [6] R. W. Beard, T. W. McLain, M. A. Goodrich, and E. P. Anderson. Coordinated target assignment and intercept for unmanned air vehicles. *IEEE Trans. Robotics and Automation*, 18(6):911–922, 2002.
- [7] P. Toth and D. Vigo. *The vehicle routing problem*. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, PA, 2002.
- [8] Li Wei and C. G. Cassandras. A cooperative receding horizon controller for multivehicle uncertain environments. *IEEE Trans. on Automatic Control*, 51(2):242–257, 2006.
- [9] 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.
- [10] D. J. Bertsimas and G. J. van Ryzin. A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Operations Research*, 39:601–615, 1991.
- [11] 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.
- [12] 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.

- [13] S. Bspamyatnikh, D. Kirkpatrick, and J. Snoeyink. Generalizing ham sandwich cuts to equitable subdivisions. *Discrete and Computational Geometry*, 24:605–622, 2000.
- [14] T. Sakai. Balanced convex partitions of measures in \mathbb{R}^2 . *Graphs and Combinatorics*, 18:169–192, 2002.
- [15] 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.
- [16] 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.
- [17] R. C. Larson and A. R. Odoni. *Urban Operations Research*. Prentice-Hall, Englewood Cliffs, NJ, 1981.
- [18] D. Applegate, R. Bixby, V. Chvátal, and W. Cook. On the solution of traveling salesman problems. In *Documenta Mathematica, Journal der Deutschen Mathematiker-Vereinigung*, pages 645–656, Berlin, Germany, August 1998. Proceedings of the International Congress of Mathematicians, Extra Volume ICM III.
- [19] 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.
- [20] S. Arora. Nearly linear time approximation scheme for Euclidean TSP and other geometric problems. In *Proc. 38th IEEE Annual Symposium on Foundations of Computer Science*, pages 554–563, Miami Beach, FL, October 1997.
- [21] S. Lin and B. W. Kernighan. An effective heuristic algorithm for the traveling-salesman problem. *Operations Research*, 21:498–516, 1973.
- [22] P. K. Agarwal and M. Sharir. Efficient algorithms for geometric optimization. *ACM Computing Surveys*, 30(4):412–458, 1998.
- [23] Z. Drezner, editor. *Facility Location: A Survey of Applications and Methods*. Springer Series in Operations Research. Springer Verlag, New York, 1995.
- [24] N. Megiddo and K. J. Supowit. On the complexity of some common geometric location problems. *SIAM Journal on Computing*, 13(1):182–196, 1984.
- [25] R. C. Larson and A. R. Odoni. *Urban Operations Research*. Prentice Hall, 1981.
- [26] H. Xu. *Optimal policies for stochastic and dynamic vehicle routing problems*. Dept. of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA., 1995.
- [27] 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.
- [28] A. Arsie, K. Savla, and E. Frazzoli. Efficient routing algorithms for multiple vehicles with no explicit communications. *IEEE Trans. Automatic Control*, 54(10):2302–2317, 2009.
- [29] M. Pavone, N. Bisnik, E. Frazzoli, and V. Isler. A stochastic and dynamic vehicle routing problem with time windows and customer impatience. *ACM/Springer Journal of Mobile Networks and Applications*, 14(3):350–364, 2009.
- [30] S. L. Smith, M. Pavone, F. Bullo, and E. Frazzoli. Dynamic vehicle routing with priority classes of stochastic demands. *SIAM Journal on Control and Optimization*, 48(5):3224–3245, 2010.
- [31] R.W. Wolff. *Stochastic modeling and the theory of queues*. Prentice Hall, 1989.
- [32] 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.
- [33] M. Pavone, E. Frazzoli, and F. Bullo. Distributed policies for equitable partitioning: Theory and applications. In *Proc. IEEE Conf. on Decision and Control*, pages 4191–4197, Cancun, Mexico, December 2008.
- [34] M. Pavone, A. Arsie, E. Frazzoli, and F. Bullo. Equitable partitioning policies for robotic networks. In *Proc. IEEE Conf. on Robotics and Automation*, pages 3979–3984, Kobe, Japan, May 2009.
- [35] M. Pavone, A. Arsie, E. Frazzoli, and F. Bullo. Equitable partitioning policies for mobile robotic networks. *IEEE Trans. Automatic Control*, August 2009. Submitted Dec. 2008, provisionally accepted. Extended version, available at <http://arxiv.org/abs/0903.5267>.
- [36] 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.
- [37] M. Pavone and E. Frazzoli. Dynamic vehicle routing with stochastic time constraints. In *Proc. IEEE Conf. on Robotics and Automation*, Anchorage, Alaska, May 2010. To Appear.