

On Dynamic Vehicle Routing with Time Constraints

Shaunak D. Bopardikar Stephen L. Smith Francesco Bullo

Abstract—We consider the problem of dynamic vehicle routing under exact time constraints on servicing demands. Demands for service are generated in an environment as follows: uniformly randomly in space and Poisson in time. Every demand needs to be serviced exactly after a fixed, finite interval of time after it is generated. We design routing policies for a service vehicle with simple kinematics to maximize the fraction of demands serviced at steady-state. The main contributions are as follows. First, we demonstrate that this problem is described by an appropriate directed acyclic graph structure which leads to a computationally-efficient routing algorithm based on a longest-path computation. Second, we provide two analytic lower bounds on the service fraction of the longest path policy. The first bound is relative to an optimal, non-causal version of the policy, i.e., a policy based on knowledge of all future demand requests. The second bound is an explicit function of demand generation rate, and therefore, useful as a design tool. We also present numerical results to support the analytic bounds and to shed light on parameter regimes where the analytic bounds are not conclusive.

I. INTRODUCTION

Dynamic Vehicle Routing (DVR) refers to a class of path planning problems for one (or many) vehicle(s) to efficiently service demand requests that appear sequentially in a given environment as per a spatio-temporal process. These problems find application in the areas of *surveillance/reconnaissance*, where the goal is to track mobile targets [1]; in *environmental monitoring*, where a dynamically evolving map on a region needs to be estimated [2]; as well as in *industrial automation*, wherein robotic arms need to perform efficient pick-and-place operations [3].

Early results on DVR problems comprised of policies that achieved the minimum (for arrival rates tending to zero), or were within a constant factor of optimality (for arrival rates tending to infinity) with respect to the expected time spent by each demand before being served [4], [5]. A single policy was proposed in [6] which is optimal for the case of low arrival rate and performs within a constant factor of the best known policy for the case of high arrival rate. Due to a recent surge of activity in the area of motion planning for autonomous robots, there have been a lot of variants of DVR being addressed over the last decade. We refer the reader to [7] for a comprehensive survey on this topic.

Pertaining to the variant of DVR problems with time windows, the problem traces its origin to the classic static vehicle routing problem with time windows [8], which is known to be NP-hard. A dynamic version of this problem was considered in [9], which also accounted for demands stochastically disappearing with a known distribution. The work in [10] considers a related problem where demands appear and disappear via known time distributions and take place at fixed points of interest in a region. Related problems in which the goal is to efficiently plan collision-free paths through environments with obstacles have been considered more recently in [11] and in [12]. Both of these two references deal with a problem which is essentially a dual of the one we consider in the sense that our goal is to reach certain points in a region.

This material is based upon work supported in part by NSF Award CPS-1035917 and ARO Award W911NF-11-1-0092.

S. D. Bopardikar is with United Technologies Research Center Inc, Berkeley, CA. Email: bshaunak@gmail.com. S. L. Smith is with the Department of Electrical and Computer Engineering, University of Waterloo, Canada. Email: stephen.smith@uwaterloo.ca. F. Bullo is with the Center for Control, Dynamical Systems and Computation, University of California at Santa Barbara, Santa Barbara, CA 93106. Email: bullo@engineering.ucsb.edu.

This paper considers a DVR problem in which demands appear uniformly randomly in a compact planar environment, assumed to be a square of length W for ease of presentation, and via a Poisson process in time with parameter λ . Each demand needs to be serviced exactly after a time duration of T after its generation. These scenarios arise in several robotic applications such as event monitoring in surveillance, or in transportation networks, when service is required at exact time instants. A demand is *missed* if it is not serviced at the specified time instant. A demand gets serviced by a vehicle, modeled as a first-order integrator with maximum speed u , when the vehicle reaches the demand location at the exact instant of time when it is to be serviced. Our goal is to design routing policies, i.e., the order in which the demands should be serviced, so as to maximize the fraction of demands serviced at steady state.

Our main contributions are as follows. First, we demonstrate that this problem is described by an appropriate directed acyclic graph structure in the space-time environment. This structure leads to a computationally-efficient routing algorithm based on a longest-path computation. Second, we provide two novel analytic lower bounds on the service fraction of the longest path policy. The first bound is a function of W, u and T and is defined relative to an optimal, non-causal version of the policy, i.e., a policy based on knowledge of all future demand requests. Such comparison with non-causal policies have been studied in literature under the terminology of *competitive ratio*, e.g., see [13] on characterizing this ratio for the k-server problem. Through this bound, we establish asymptotic optimality of the longest path policy in the parameter regimes for which the term $W/(Tu) \rightarrow 0^+$. However, this bound is not amenable for use as a design tool primarily because the non-causal version of the policy is not physically realizable. For the case when $T \geq \sqrt{2}W/u$, we derive a second bound, which is an explicit function of W, u and λ , and therefore, useful as a design tool. Finally, we present numerical results to support the analytic bounds as well as to shed light on parameter regimes where the analytic bounds are not conclusive.

A preliminary version of this work was addressed in [14], in which we considered the case of mobile (translating) demands being generated on a line segment, and which need to be serviced before they reached a finish line. While the present work borrows the main concepts, such as identifying the directed acyclic graph structure, the novelty of this paper lies in formulating the problem in space-time environment, and in the derivation of the two analytic lower bounds.

This paper is organized as follows. The problem formulation along with background results are presented in Section II. The service policies are described in Section III. The analytic lower bounds are presented in Section IV. Simulation results are presented in Section V. Finally, conclusions and directions for future work are presented in Section VI.

II. PROBLEM STATEMENT AND BACKGROUND

In this section, we present the problem statement and background results useful to establish the main results of this paper.

A. Problem Formulation

Consider a square environment $\mathcal{E} := [0, W]^2 \subset \mathbb{R}^2$. The environment contains a single vehicle with position $\mathbf{p}(t) = [X(t), Y(t)]^T \in \mathcal{E}$, modeled as a first-order integrator with maximum speed equal to u . Demands for service are generated in the environment via a spatio-temporal process. We assume that the process generating the demands is *uniform in space* and *Poisson in time* with parameter λ . Specifically, we assume that if a tagged demand $i, \forall i \in \mathbb{N}$, is generated at time $t_{\text{rel},i}$, then it is required to be serviced at the exact time instant given by $t_{\text{rel},i} + T$, for a given $T > 0$. The demand is *served* when the

vehicle reaches exactly the spatial location of the demand at the time instant $t_{\text{rel},i} + T$.

Remark II.1 (Finite service time windows) The algorithms in this paper and the analysis also extend to the case when a tagged demand i , for all $i \in \mathbb{N}$, is generated at time $t_{\text{rel},i}$, needs to be serviced at any time in the interval $[t_{\text{rel},i}, t_{\text{rel},i} + T]$. However, the analysis would become conservative for the case, especially when the generation rate λ is not very high.

Let $\mathcal{Q}(t) \subset \mathcal{E}$ denote the set of positions of all released but unserved demands at time t . If the i th demand is served, then it is removed from \mathcal{Q} and placed in the set \mathcal{Q}_{srv} , whose cardinality we denote m_{srv} . If the i th demand is missed, then it is removed from \mathcal{Q} and placed in $\mathcal{Q}_{\text{miss}}$ with cardinality m_{miss} .

Online and Offline Algorithms: An online algorithm [15] (or policy) for the vehicle is a map $\mathcal{P} : \mathcal{E} \times \mathbb{F}(\mathcal{E}) \rightarrow \mathbb{R}^2$, where $\mathbb{F}(\mathcal{E})$ is the set of finite subsets of \mathcal{E} , assigning a commanded velocity to the service vehicle as a function of the current state of the system: $\dot{\mathbf{p}}(t) = \mathcal{P}(\mathbf{p}(t), \mathcal{Q}(t))$. Thus, in an online algorithm the vehicle at a time t has access to demand information, only for the demands that have been generated until time t . By contrast, in an offline algorithm, the vehicle has access to the generation time and the location of all demands throughout its entire execution. Thus, the vehicle trajectory $t \mapsto \mathbf{p}(t)$ can be computed at the problem outset. In particular, if $\bar{\mathcal{Q}}$ denotes the set of positions of all demands that will be released in the execution, then an offline algorithm can be analogously described by the form $\dot{\mathbf{p}}(t) = \mathcal{P}(\mathbf{p}(t), \bar{\mathcal{Q}})$.

Problem Statement: The goal in this paper is to find online algorithms P that maximize the fraction of demands that are serviced $\mathbb{F}_{\text{cap}}(P)$, termed as the *service fraction*. Formally, for a policy P , we define the steady state average service fraction as

$$\mathbb{F}_{\text{cap}}(P) := \limsup_{t \rightarrow +\infty} \mathbb{E} \left[\frac{m_{\text{srv}}(t)}{m_{\text{srv}}(t) + m_{\text{miss}}(t)} \right],$$

where the expectation is with respect to the stochastic process that generates the demands.

B. Background Results

We now review the distribution of demands in an unserved region and the longest path problem.

1) *Demand distribution:* Suppose that mobile demands are generated uniformly randomly in the region \mathcal{E} and by a Poisson process in time with rate λ . Upon generation, each demand moves in the $+z$ direction with unit speed, i.e., all demands translate in the $+z$ direction. Then, the following result characterizes the distribution of demands within any region contained in $\mathcal{E} \times [0, t]$.

Lemma II.2 (Distribution of outstanding demands, [16])

Suppose the generation of demands commences at time 0 and no demands are serviced in the interval $[0, t]$. Let \mathcal{Q} denote the set of all demands in $\mathcal{E} \times [0, t]$ at time t . Then, given a measurable compact region \mathcal{R} of volume V contained in $\mathcal{E} \times [0, t]$,

$$\mathbb{P}[|\mathcal{R} \cap \mathcal{Q}| = n] = \frac{e^{-\lambda V/W^2} (\lambda V/W^2)^n}{n!}.$$

As a consequence, conditioned on the number of demands within such a region \mathcal{R} , the demands are distributed uniformly randomly in \mathcal{R} .

The proof of this result is identical to the proof of [16][Lemma V.1] in which we substitute the demand speed $v = 1$, and therefore, we omit the proof for the sake of brevity.

2) *Longest Paths in Directed Acyclic Graphs:* A directed graph $G = (V, E)$ consists of a set of vertices V and a set of directed edges $E \subset V \times V$. An edge $(v, w) \in E$ is directed from vertex v to vertex w . A *path* in G is a sequence of vertices such that from each vertex in the sequence, there is an edge in E directed to the next vertex in the sequence. A path is *simple* if it contains no repeated vertices. A *cycle* is a path in which the first and last vertices in the sequence are the same. A graph G is *acyclic* if it contains no cycles. The longest path problem is to find a simple path of maximum length (i.e., a path that visits a maximum number of vertices). In general this problem is NP-hard as its solution would imply a solution to the well known Hamiltonian path problem [17]. However, if the graph is a Directed Acyclic Graph (DAG), then the longest path problem has an efficient dynamic programming solution [18] with complexity $O(|V| + |E|)$, that relies on topologically sorting [19] the vertices.

III. SERVICE POLICIES

In this section, we begin by introducing a notion of reachability graph, and then proceed to define service policies based on the computation of longest paths over the reachability graph.

A. Characterization of reachable demands

Without loss of generality, we assume that the demands are labelled sequentially as per their generation time. When the i th demand is active, we can attach a third-coordinate to the position of the demand, given by

$$t_{\text{rel},i} + T - t,$$

where $t_{\text{rel},i}$ is the generation time of the demand, and t is the current time. Therefore, we define the space-time environment as $\mathcal{E}_{\text{ST}} := \mathcal{E} \times \mathbb{R}_{\geq 0}$.

Consider demand i with generation time $t_{\text{rel},i}$ and position $\mathbf{q}_i = (x_i, y_i) \in \mathcal{E}$. Suppose that the service vehicle is located at $\mathbf{p}(t) = (X(t), Y(t))$, at a time instant t satisfying $t_{\text{rel},i} \leq t \leq t_{\text{rel},i} + T$. The demand i can be serviced if and only if

$$\frac{\|\mathbf{p}(t) - \mathbf{q}_i\|}{u} \leq t_{\text{rel},i} + T - t.$$

The left hand side is the time it takes for the vehicle to reach \mathbf{q}_i from its current position $\mathbf{p}(t)$ moving at full speed toward the demand location, while the right hand side is the remaining time until the demand's service constraint. From this description, we can define the set of reachable demands from any location in \mathcal{E}_{ST} .

Definition III.1 (Reachable set from a point) The *reachable set* $R_T(\mathbf{y}, t)$ from a position $\mathbf{y} \in \mathcal{E}$ at time $t \geq 0$ is defined as

$$R_T(\mathbf{y}, t) := \{(\mathbf{z}, \tau) \in \mathcal{E} \times [t, T] \mid \frac{\|\mathbf{y} - \mathbf{z}\|}{u} \leq \tau - t\}.$$

Next, consider the set of demands in $R_T(\mathbf{p}(\bar{t}), \bar{t})$, and suppose the vehicle chooses to service demand i , with position $(\mathbf{q}_i, t_i) = (x_i, y_i, t_i) \in R_T(\mathbf{p}(\bar{t}), \bar{t})$. Upon service at time $t_i = t_{\text{rel},i} + T$, the service vehicle can recompute the reachable set $R_T(\mathbf{p}(t_i), t_i)$, and select a demand that lies within. Since all demands translate together along the time axis, every demand that is reachable from (\mathbf{q}_i, t_i) , is reachable from (\mathbf{q}_i, \bar{t}) . Thus, the service vehicle can “look ahead in time with a horizon of T ” and compute the demands that will be reachable from each serviced demand position in \mathcal{E}_{ST} . This idea motivates the concept of a reachability graph in \mathcal{E}_{ST} .

Definition III.2 (Reachability graph) The *reachability graph* of a set of points $\{(\mathbf{q}_1, t_1), \dots, (\mathbf{q}_n, t_n)\} \in \mathcal{E}_{\text{ST}}$, is a directed acyclic graph with vertex set $V := \{1, \dots, n\}$, and edge set E , where for

$i, j \in V$, the edge (i, j) is in E if and only if $(\mathbf{q}_j, t_j) \in R_T(\mathbf{q}_i, t_i)$ and $j \neq i$.

Given a set Q of n outstanding demands, and a vehicle position (\mathbf{p}, t) , we can compute the corresponding reachability graph (see Fig. 1). In addition, by Section II-B2, we can compute the longest

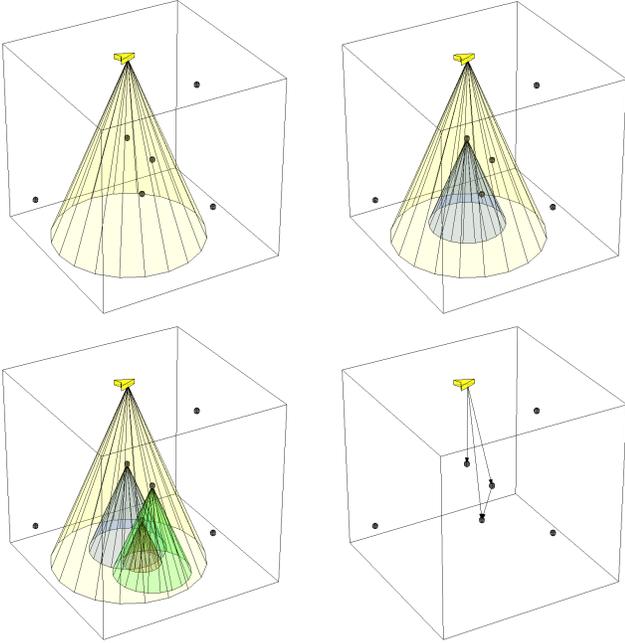


Fig. 1. The construction of the reachability graph. The top-left figure shows the set of reachable demand locations in the space-time environment \mathcal{E}_{ST} from the vehicle positioned at a location in the environment \mathcal{E} . The top-right and bottom-left figures show the reachable set R_T from the applicable demand locations in \mathcal{E}_{ST} . The bottom-right figure shows the reachability graph in \mathcal{E}_{ST} .

path in a reachability graph in $O(n^2)$ computation time.

B. The Longest Path Policy

We now introduce the Longest Path policy. In the LP policy, the fraction η is a design parameter. The lower η is chosen, the better the performance of the policy, but this comes at the expense of increased computation.

Algorithm 1: The Longest Path (LP) policy

- 1 Compute the reachability graph of the vehicle position and all unserved demands in $Q(0)$.
 - 2 Compute a longest path in this graph, starting at the service vehicle location.
 - 3 **if** longest path is empty **then**
 - 4 Move to the center of \mathcal{E} and recompute reachability graph whenever a new demand is released, and return to step 2.
 - 5 **else**
 - 6 Service demands in the order they appear on the path.
 - 7 Once a fraction $\eta \in]0, 1]$ of the demands on the path have been serviced, recompute the reachability graph of all outstanding demands and return to step 2.
-

C. A Non-causal Longest Path Policy

For the sake of characterizing the performance of the Longest Path policy, we will consider a non-causal policy. In the online algorithms

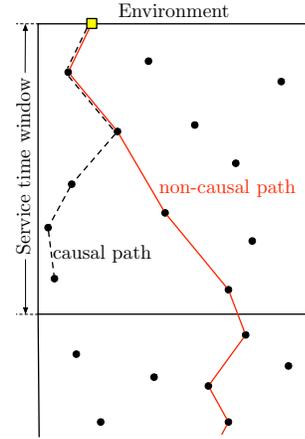


Fig. 2. A snapshot in the evolution of the Non-causal Longest Path policy as shown in the \mathcal{E}_{ST} viewed along the $+y$ direction. The vehicle has planned the solid red path through all demands, including those that have not yet arrived. In comparison, a dashed causal longest path is shown, which only considers demands that have arrived.

literature, such a policy is referred to as an *offline algorithm* [15]. Fig. 2 shows an example of a path generated by the Non-causal Longest Path policy. Note that the service vehicle will serve each demand by moving to its location in \mathcal{E} , and thus the path depicts which demands will be served, and in what order.

Algorithm 2: Non-causal Longest Path (NCLP) policy

Assumes: Release times for all demands are known a priori.

- 1 Compute the reachability graph of the vehicle position and all demands in $Q(0) \cup \bar{Q}(0)$.
 - 2 Compute a longest path in this graph, starting at the service vehicle location.
 - 3 Serve demands in the order they appear on the path.
-

While this policy is not physically realizable, it will serve as a baseline to evaluate the performance of a causal, longest path policy as described in the next section.

IV. ANALYSIS OF THE LONGEST PATH POLICY

In this section, we present analytic results characterizing bounds on the performance of the Longest Path Policy, defined in Section III. The LP policy is difficult to analyze directly. This is due to the fact that the position of the vehicle at time t depends on the positions of all outstanding demands in $Q(t)$. Therefore, in this section, we will present two approaches to analyze the policy; the first is to derive a factor of optimality with respect to the Non-causal Longest Path policy, and the second is to lower bound the performance of the LP policy by comparing it with a simpler, greedy policy.

A. Comparison with Non-causal Longest Path

The first approach is summarized by the following theorem in which we relate the Longest Path policy to its non-causal relative. Such a bound is referred to as a *competitive ratio* in the online algorithms literature [15].

Theorem IV.1 (Optimality of Longest Path policy) *The service fraction for the Longest Path Policy satisfies*

$$\mathbb{F}_{\text{cap}}(\text{LP}) \geq \left(1 - \frac{\sqrt{2}W}{Tu}\right) \mathbb{F}_{\text{cap}}(\text{NCLP}),$$

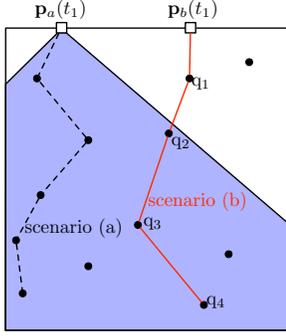


Fig. 3. Scenario (a) and (b) for the proof of Theorem IV.1, as viewed along the $+y$ direction. Path (a) visits five demands and thus $\mathcal{L}_a = 5$. Path (b) visits four demands, yielding $m = 4$. The demand \mathbf{q}_2 is the highest on path (b) that can be serviced from $\mathbf{p}_a(t_1)$. Thus, $n = 1$, and $5 = \mathcal{L}_a > m - n = 3$.

so that the Longest Path policy is optimal as $W/Tu \rightarrow 0^+$.

Proof: Suppose that the generation of demands begins at $t = 0$ and let us consider two scenarios; (a) the vehicle uses the Longest Path policy, and (b) the vehicle uses the Non-causal Longest Path policy. Then, at any instant in time $t_1 > 0$ we can compare the number of demands serviced in scenario (a) to the number serviced in scenario (b) (refer to Fig. 3).

Let us consider a time instant t_1 where in scenario (a), the vehicle is recomputing the longest path through all outstanding demands $\mathcal{Q}(t_1)$. Let $\mathbf{p}_a(t_1)$ and $\mathbf{p}_b(t_1)$ denote the vehicle position in scenario (a) and scenario (b), respectively, at time t_1 . In scenario (b), let the path that the vehicle will take through $\mathcal{Q}(t_1)$ be given by $(\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_m) \in \mathcal{Q}(t_1)$. The demand \mathbf{q}_1 is reachable from $\mathbf{p}_b(t_1)$, but it may not be reachable from $\mathbf{p}_a(t_1)$. However, a lower bound on the length of the longest path in scenario (a) is: $(\mathbf{q}_{n+1}, \mathbf{q}_{n+2}, \dots, \mathbf{q}_m)$, where \mathbf{q}_{n+1} , $n \in \{0, \dots, m-1\}$, is the demand with the smallest release time that can be reached from $\mathbf{p}_a(t_1)$, see Fig. 3. Thus, the length of the longest path in scenario (a), \mathcal{L}_a , is at least

$$\mathcal{L}_a \geq m - n, \quad (1)$$

where m is the length of the path in scenario (b).

Now, since the environment has diameter $\sqrt{2}W$, the vehicle in scenario (a) can service any demand $(\mathbf{q}_i, t_{\text{rel},i})$ with $t_{\text{rel},i} + T - t_1 \geq \sqrt{2}W/u$. Thus, the demands $\mathbf{q}_1, \dots, \mathbf{q}_n$ must all have their release times in $]t_1 - T, t_1 + \sqrt{2}W/u - T]$, since they are active at time t_1 . Let the total number of active outstanding demands at time t_1 be N_{tot} . Then, conditioned on N_{tot} , by Lemma II.2, the expected number of active demands contained in $[0, W] \times [0, W] \times]t_1 - T, t_1 + \sqrt{2}W/u - T]$ is $N_{\text{tot}}\sqrt{2}W/(Tu)$. Hence,

$$\mathbb{E}[n|N_{\text{tot}}] = N_{\text{tot}} \frac{\sqrt{2}W}{Tu} \mathbb{F}_{\text{cap}}(\text{NCLP}). \quad (2)$$

Similarly, for the length of the path through $\mathcal{Q}(t_1)$ in scenario (b), we have

$$\mathbb{E}[m|N_{\text{tot}}] = N_{\text{tot}} \mathbb{F}_{\text{cap}}(\text{NCLP}). \quad (3)$$

Combining equations (2) and (3) with equation (1) we obtain

$$\begin{aligned} \mathbb{E}[\mathcal{L}_a|N_{\text{tot}}] &\geq N_{\text{tot}} \left(1 - \frac{\sqrt{2}W}{Tu}\right) \mathbb{F}_{\text{cap}}(\text{NCLP}), \\ \mathbb{E}\left[\frac{\mathcal{L}_a}{N_{\text{tot}}}|N_{\text{tot}}\right] &\geq \left(1 - \frac{\sqrt{2}W}{Tu}\right) \mathbb{F}_{\text{cap}}(\text{NCLP}). \end{aligned}$$

But $\mathcal{L}_a/N_{\text{tot}}$ is the fraction of outstanding demands in $\mathcal{Q}(t_1)$ that will be serviced in scenario (a), and it does not depend on the value

of N_{tot} . By the law of total expectation

$$\mathbb{E}\left[\frac{\mathcal{L}_a}{N_{\text{tot}}}\right] = \mathbb{E}\left[\mathbb{E}\left[\frac{\mathcal{L}_a}{N_{\text{tot}}}|N_{\text{tot}}\right]\right] \geq \left(1 - \frac{\sqrt{2}W}{Tu}\right) \mathbb{F}_{\text{cap}}(\text{NCLP}).$$

At each epoch when the longest path is recomputed, the path in scenario (a) will service at least this fraction of demands. Thus, we have $\mathbb{F}_{\text{cap}}(\text{LP}) \geq \mathbb{E}[\mathcal{L}_a/N_{\text{tot}}]$ and have proved the result. ■

Note that in this approach, the performance has been compared to the optimal non-causal version. Theorem IV.1 does not provide a bound on the performance of the LP policy in an absolute sense, e.g., as a function of the environment dimensions or the demand generation rate. Such a bound requires a different analysis, presented in the following subsection.

B. Lower bound using a Greedy policy

The second approach to characterize the LP policy performance is to lower bound the service fraction of the LP policy with a greedy policy. We first define the notion of an *inner reachability set*, which essentially serves as an inner approximation to the reachability set from a location.

Definition IV.2 (Inner Reachability Set) Given a vehicle location $\mathbf{p}(t) = (X(t), Y(t))$, the inner reachability set $R_{\text{in},T}(\mathbf{p}, t)$ is a subset of the reachability set $R_T(\mathbf{p}, t)$ such that

$$\begin{aligned} R_{\text{in},T} = \bigcup_{\tau \in [0, T]} & \left[X - \frac{u(T-\tau)}{\sqrt{2}}, X + \frac{u(T-\tau)}{\sqrt{2}} \right] \times \\ & \left[Y - \frac{u(T-\tau)}{\sqrt{2}}, Y + \frac{u(T-\tau)}{\sqrt{2}} \right] \times \{t + T - \tau\}. \end{aligned}$$

In other words, this set is a pyramid whose base is the square region $[X - uT/\sqrt{2}, X + uT/\sqrt{2}] \times [Y - uT/\sqrt{2}, Y + uT/\sqrt{2}]$. Since this set is contained inside $R_T(X, Y, t)$, any demand inside $R_{\text{in},T}$ can be served by the vehicle. We will now define a Greedy Path policy.

Algorithm 3: The Greedy Path (GP) policy

- Assumes:** Vehicle is located at $(X(t), Y(t))$
- 1 Compute the inner reachability set $R_{\text{in},T}(X, Y, t)$.
 - 2 Service the demand in $R_{\text{in},T}(X, Y, t)$ with the highest time coordinate.
 - 3 If no demands exist, move toward the center of the environment.
 - 4 Repeat.
-

Given a set of outstanding demands $\mathcal{Q}(t)$ at time t , the vehicle position is independent of all outstanding demands, except the demand currently being serviced. At any time instant, let $\ell, g \in \mathcal{E}$ denote the positions of the vehicle when following the Longest path and the Greedy path policies from the initial instant, respectively. The capture fraction \mathbb{F}_{cap} can be viewed as a function of \mathcal{Q} and the vehicle position along with the choice of policy P . Since both ℓ and g are independent of \mathcal{Q} ,

$$\mathbb{F}_{\text{cap}}(\text{LP}, \mathcal{Q}, \ell) = \mathbb{F}_{\text{cap}}(\text{LP}, \mathcal{Q}, g).$$

But conditioned on a position ℓ (or g),

$$\mathbb{F}_{\text{cap}}(\text{LP}, \mathcal{Q}) \geq \mathbb{F}_{\text{cap}}(G, \mathcal{Q}),$$

since the Greedy Path policy generates a suboptimal longest path through $\mathcal{Q}(t)$. Combining these two relations, we conclude that

$$\mathbb{F}_{\text{cap}}(\text{LP}, \mathcal{Q}) \geq \mathbb{F}_{\text{cap}}(G, \mathcal{Q}),$$

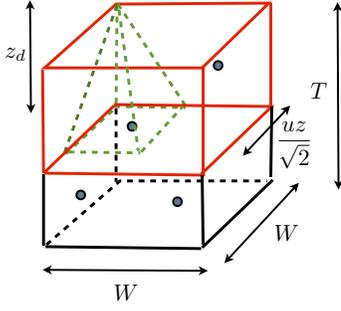


Fig. 4. The setup for the proof of Theorem IV.3. The service vehicle is located at $(0, 0, t)$. The worst case for the region $R_{\text{in}, T, z}$ is indicated by the green dashed line. As per the Greedy policy, any demand in the red region, which is outside of the dashed green region is left unserved.

which implies that the Greedy Path policy provides a lower bound on the performance of the Longest path policy.

We are now able to establish the following result.

Theorem IV.3 (Lower Bound for Longest Path policy) *If $T \geq \sqrt{2}W/u$, then for the Longest Path policy*

$$\mathbb{F}_{\text{cap}}(\text{LP}) \geq \mathbb{F}_{\text{cap}}(\text{GP}) \geq 1/g(\lambda, u, W),$$

where the function

$$g(\lambda, u, W) := e^{-\frac{\sqrt{2}\lambda W}{3u}} + \frac{\lambda}{3} \left(\frac{6W^2}{\lambda u^2} \right)^{\frac{1}{3}} \left(\Gamma\left(\frac{1}{3}, 0\right) - \Gamma\left(\frac{1}{3}, \frac{\sqrt{2}\lambda W}{3u}\right) \right),$$

and the function

$$\Gamma(a, z) := \int_z^{+\infty} t^{a-1} e^{-t} dt$$

is the incomplete Gamma function.

Proof: We begin by looking at the expression for the service fraction. Notice that if $m_{\text{srv}}(t) > 0$ for some $t > 0$, then

$$\begin{aligned} \limsup_{t \rightarrow +\infty} \mathbb{E} \left[\frac{m_{\text{srv}}(t)}{n_{\text{srv}}(t) + n_{\text{miss}}(t)} \right] &= \limsup_{t \rightarrow +\infty} \mathbb{E} \left[\frac{1}{1 + \frac{m_{\text{miss}}(t)}{m_{\text{srv}}(t)}} \right] \\ &\geq \left(1 + \limsup_{t \rightarrow +\infty} \mathbb{E} \left[\frac{m_{\text{miss}}(t)}{m_{\text{srv}}(t)} \right] \right)^{-1}, \end{aligned} \quad (4)$$

where the last step comes from an application of Jensen's inequality [20]. Thus, we can determine a lower bound on the service fraction by studying the number of demands that escape per serviced demand.

Let us study the time instant t at which the service vehicle services its i th demand, and determine an upper bound on the number of demands that escape before the service vehicle services its $(i+1)$ th demand. Since we seek a lower bound on the service fraction of the LP policy, we may consider the path generated by the Greedy Path policy. In addition, we consider the worst-case service vehicle position in \mathcal{E}_{ST} ; namely, the position $(0, 0, t)$ (or equivalently, any of the top four corners in Fig. 4).

From any position of the vehicle, the reachable set R_T is a cone whose axis is oriented in the negative z direction, and the half-angle of this cone is numerically equal to $\arctan(u)$. Let $R_{\text{in}, T, z}$ denote the inner reachability set intersected with $[0, W] \times [0, W] \times [t+T-z, t+T]$, where $z \in [0, T]$, and let $|R_{\text{in}, T, z}|$ denote its volume. Then,

$$|R_{\text{in}, T, z}| = \begin{cases} \frac{u^2 z^3}{6}, & \text{if } z \leq \sqrt{2}W/u, \\ \frac{\sqrt{2}\lambda W^3}{3u} + W^2(z - \frac{\sqrt{2}W}{u}), & \text{if } z > \sqrt{2}W/u, \end{cases}$$

An illustration of the set $R_{\text{in}, T, z}$ is shown in Fig. 4. Let z_d be the

z -distance to the reachable demand with the highest z -coordinate. That is,

$$z_d := \min_{(x, y, z) \in \mathcal{Q}(t) \cap R_T(0, 0, t)} \{T - z\},$$

where $\mathcal{Q}(t)$ is the set of outstanding demands at time t . By Lemma II.2, the probability that a subset $\mathcal{B} \subset \mathcal{E}$ with volume $|\mathcal{B}|$ contains zero demands is given by

$$\mathbb{P}[|\mathcal{B} \cap \mathcal{Q}(t)| = 0] = e^{-\lambda|\mathcal{B}|/W^2},$$

where $|\mathcal{B} \cap \mathcal{Q}(t)|$ denotes the cardinality of the finite set $\mathcal{B} \cap \mathcal{Q}(t)$. Thus,

$$\mathbb{P}[z_d > z] = \mathbb{P}[|R_{\text{in}, T, z} \cap \mathcal{Q}(t)| = 0] = e^{-\lambda|R_{\text{in}, T, z}|/W^2}.$$

The probability density function of z_d for $z \leq \sqrt{2}W/u$, which is obtained from the expression

$$f(z) = \frac{d}{dz} (1 - \mathbb{P}[z_d > z]) = -\frac{d}{dz} e^{-\lambda|R_{\text{in}, T, z}|/W^2},$$

satisfies

$$f(z) = \frac{\lambda u^2}{2W^2} z^2 e^{-\lambda u^2 z^3 / (6W^2)}.$$

Now, given z_d , all demands residing in the region $\text{miss}_{z_d} := ([0, W] \times [0, W] \times [T - z_d, T]) \setminus R_{\text{in}, T, z_d}$ will be lost unserved (see Fig. 4). Using the expressions for $|R_{\text{in}, T, z}|$, the volume of miss_{z_d} is given by

$$|\text{miss}_{z_d}| = \begin{cases} z_d W^2 - \frac{u^2 z_d^3}{6}, & \text{if } z_d \leq \sqrt{2}W/u, \\ 2\sqrt{2}\lambda W^3 / (3u), & \text{if } z_d \geq \sqrt{2}W/u. \end{cases}$$

From Lemma II.2, the expected number of outstanding demands in an unserved region of volume V is $\lambda V/W^2$. Thus, given that the vehicle is located at $(0, 0, t) \in \mathcal{E}_{\text{ST}}$, the expected number of demands that will be missed while the service vehicle is serving its $(i+1)$ th demand is given by

$$\begin{aligned} \mathbb{E}[m_{\text{miss}, i}] &= \frac{\lambda}{W^2} \mathbb{E}[|\text{miss}_{z_d}|] \\ &= \frac{\lambda}{W^2} \left[\int_0^{\sqrt{2}W/u} \left(zW^2 - \frac{u^2 z^3}{6} \right) f(z) dz \right. \\ &\quad \left. + \frac{2\sqrt{2}\lambda W^3}{3u} \mathbb{P}[z_d > \sqrt{2}W/u] \right]. \end{aligned}$$

Applying the probability density function and cumulative distribution function of z_d we obtain

$$\begin{aligned} \mathbb{E}[m_{\text{miss}, i}] &= \frac{\lambda^2 u^2}{2W^4} \int_0^{\sqrt{2}W/u} \left(W^2 z^3 - \frac{u^2 z^5}{6} \right) e^{-\lambda u^2 z^3 / (6W^2)} dz \\ &\quad + \frac{2\sqrt{2}\lambda W}{3u} e^{-\sqrt{2}\lambda W / (3u)}. \end{aligned} \quad (5)$$

The integral can be evaluated using Mathematica¹ which gives

$$\begin{aligned} \mathbb{E}[m_{\text{miss}, i}] &= \left(1 - \frac{2\sqrt{2}\lambda W}{3u} \right) e^{-\frac{\sqrt{2}\lambda W}{3u}} + \frac{2\sqrt{2}\lambda W}{3u} e^{-\frac{\sqrt{2}\lambda W}{3u}} \\ &\quad + \frac{\lambda}{3} \left(\frac{6W^2}{\lambda u^2} \right)^{\frac{1}{3}} \left(\Gamma\left(\frac{1}{3}, 0\right) - \Gamma\left(\frac{1}{3}, \frac{\sqrt{2}\lambda W}{3u}\right) \right) - 1. \end{aligned}$$

Since $\mathbb{E}[m_{\text{miss}, i}]$ is computed for the worst-case vehicle position $(0, 0, t)$ in \mathcal{E}_{ST} , and since this expression holds at every demand service, we have that

$$\begin{aligned} \limsup_{t \rightarrow +\infty} \mathbb{E} \left[\frac{m_{\text{miss}}(t)}{m_{\text{srv}}(t)} \right] &\leq e^{-\frac{\sqrt{2}\lambda W}{3u}} \\ &\quad + \frac{\lambda}{3} \left(\frac{6W^2}{\lambda u^2} \right)^{\frac{1}{3}} \left(\Gamma\left(\frac{1}{3}, 0\right) - \Gamma\left(\frac{1}{3}, \frac{\sqrt{2}\lambda W}{3u}\right) \right) - 1, \end{aligned}$$

¹Using the online integrator at <http://integrals.wolfram.com>

and thus by equation (4) we obtain the desired result. ■

Remark IV.4 (Asymptotic performance and conservativeness)

Both analytic bounds lead to asymptotic optimality in the regimes when $W/(Tu) \rightarrow 0^+$ for the first bound, or when $\lambda \rightarrow 0^+$ or $W/u \rightarrow 0^+$ in the second bound. One of the sources of conservativeness in the analysis is the use of the inner reachability set $R_{in,T}$ to define the Greedy path policy. Instead, one could use the reachability set R_T to define the Greedy path policy. However, to the best of our knowledge, that complicates the integration in (5).

V. SIMULATIONS

We now present results from numerical experiments. Through these experiments, we compare the Longest Path policy with $\eta = 1$ to the Non-causal Longest Path policy and to the theoretical lower bound in Theorem IV.3, in order to verify the analytic claims.

To simulate the LP and the NCLP policies, we perform 20 runs of the policy, where each run consists of 500 demands. A comparison of the service fractions for the two policies is presented in Fig. 5. When $T > \sqrt{2}W/u$, we observe that the service fraction of the LP policy is nearly identical to that of the NCLP policy, as shown in Fig. 5(a). We also confirm the analytic results, i.e., the bound relative to the non-causal Longest path policy from Theorem IV.1, and the explicit lower bound from Theorem IV.3.

In Fig. 5(b), where $T < \sqrt{2}W/u$, the service fraction of the LP policy is observed to be within 2% of the NCLP policy. These numerical results suggest that the Longest Path policy is essentially optimal over a larger range of parameter values than what our theoretical analysis can characterize.

VI. CONCLUSIONS AND FUTURE DIRECTIONS

This paper considered a dynamic version of the classic vehicle routing problem in which service needed to be provided exactly at a specified instant of time after every demand generation. The main contribution was the identification of a directed acyclic graph structure on the problem which leads to a novel policy based on repeated computation of longest paths through the available set of demands. A performance analysis in terms of the expected fraction of demands serviced was presented, for which we provided two novel lower bounds. The first bound was relative to a non-causal version of the longest path policy, and the second was an explicit bound as a function of the problem parameters. Numerical results verifying the analytic claims were presented.

In future, it would be of interest to analyze the parameter regimes in which the bounds are not currently known to hold. Other related directions are the inclusion of dynamics in the motion of the vehicle, limited sensing range for the vehicle and multi-vehicle versions of this problem.

REFERENCES

- [1] R. Szechtman, M. Kress, K. Lin, and D. Cfor, "Models of sensor operations for border surveillance," *Naval Research Logistics*, vol. 55, no. 1, pp. 27–41, 2008.
- [2] S. L. Smith and D. Rus, "Multi-robot monitoring in dynamic environments with guaranteed currency of observations," in *IEEE Conf. on Decision and Control*, Atlanta, GA, USA, Dec. 2010, pp. 514–521.
- [3] P. Chalasani and R. Motwani, "Approximating capacitated routing and delivery problems," *SIAM Journal on Computing*, vol. 28, no. 6, pp. 2133–2149, 1999.
- [4] D. J. Bertsimas and G. J. van Ryzin, "A stochastic and dynamic vehicle routing problem in the Euclidean plane," *Operations Research*, vol. 39, no. 4, pp. 601–615, 1991.
- [5] —, "Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles," *Operations Research*, vol. 41, no. 1, pp. 60–76, 1993.

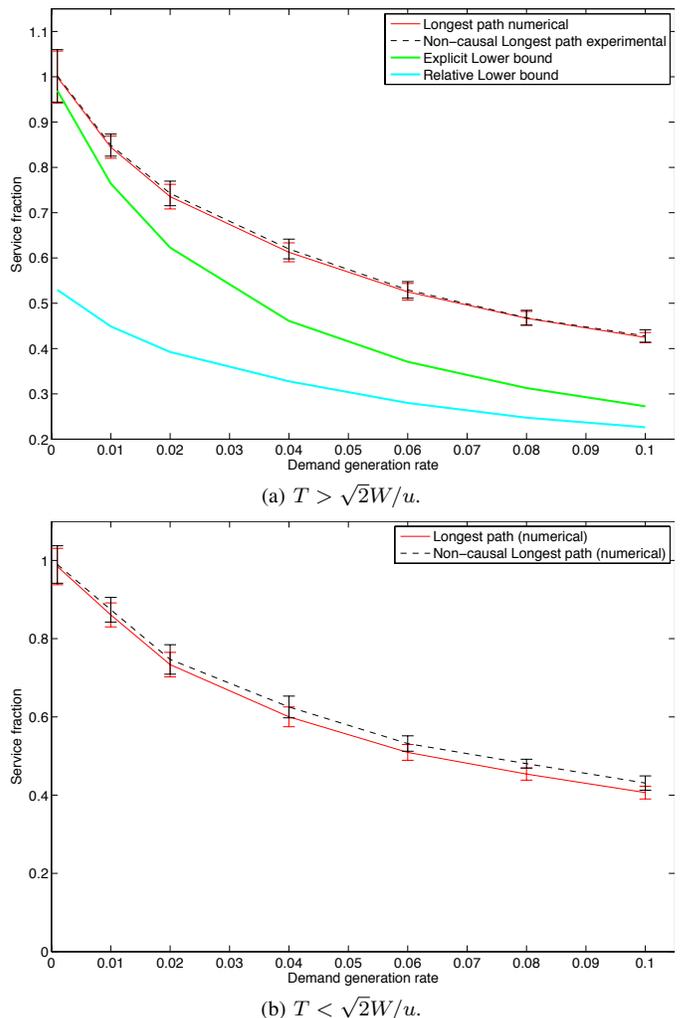


Fig. 5. Simulation results for LP policy (solid red line with error bars showing \pm one standard deviation) and the NCLP policy (dashed black line) for an environment of width $W = 100$ and $u = 3$. In (a), $T = 100 > \sqrt{2}W/u$, and the explicit lower bound in Theorem IV.3 is shown in solid green, while the lower bound relative to the non-causal Longest path is shown in solid cyan.

- [6] J. D. Papastavrou, "A stochastic and dynamic routing policy using branching processes with state dependent immigration," *European Journal of Operational Research*, vol. 95, pp. 167–177, 1996.
- [7] F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith, "Dynamic vehicle routing for robotic systems," *Proceedings of the IEEE*, vol. 99, no. 9, pp. 1482–1504, 2011.
- [8] P. Toth and D. Vigo, Eds., *The Vehicle Routing Problem*, ser. Monographs on Discrete Mathematics and Applications. SIAM, 2001.
- [9] 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*, vol. 14, no. 3, pp. 350–364, 2009.
- [10] N. Bisnik, A. A. Abouzeid, and I. Volkan, "Fast unfolding of communities in large networks," *IEEE Transactions on Robotics*, vol. 23, no. 4, pp. 676–692, 2007.
- [11] S. Karaman and E. Frazzoli, "High-speed flight in an ergodic forest," *IEEE Transactions on Robotics*, 2012, submitted. Available online at <http://sertac.scripts.mit.edu/web/wp-content/papercite-data/pdf/karaman.frazzoli-tro12.pdf>.
- [12] J. Johnson and K. Hauser, "Optimal longitudinal control planning with moving obstacles," in *2013 IEEE Int. Intelligent Vehicles Symposium*, Jun. 2013, to appear.
- [13] E. Koutsoupias and C. H. Papadimitriou, "On the k-server conjecture," *Journal of the ACM*, vol. 42, no. 5, pp. 971–983, 1995.
- [14] S. L. Smith, S. D. Bopardikar, and F. Bullo, "A dynamic boundary

- guarding problem with translating demands,” in *IEEE Conf. on Decision and Control and Chinese Control Conference*, Shanghai, China, Dec. 2009, pp. 8543–8548.
- [15] A. Borodin and R. El-Yaniv, *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
- [16] S. D. Bopardikar, S. L. Smith, F. Bullo, and J. P. Hespanha, “Dynamic vehicle routing for translating demands: Stability analysis and receding-horizon policies,” *IEEE Transactions on Automatic Control*, vol. 55, no. 11, pp. 2554–2569, 2010.
- [17] B. Korte and J. Vygen, *Combinatorial Optimization: Theory and Algorithms*, 4th ed., ser. Algorithmics and Combinatorics. Springer, 2007, vol. 21.
- [18] N. Christofides, *Graph Theory: An Algorithmic Approach*. Academic Press, 1975.
- [19] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. MIT Press, 2001.
- [20] L. Breiman, *Probability*, ser. Classics in Applied Mathematics. SIAM, 1992, vol. 7, corrected reprint of the 1968 original.