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 are sequentially generated in an environment and 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 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, under the assumption of the demands being generated uniformly randomly in the environment and via a Poisson process in time, 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 the vehicle dynamics and demand generation rate, and therefore, useful as a design tool. We present numerical results to support the analytic bounds.**

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 arise in robotic applications such as surveillance/reconnaissance, where the goal is to detect and/or track mobile targets [1], [2]; environmental monitoring, where a dynamically evolving region needs to be estimated [3]; and industrial automation, wherein robotic arms need to perform efficient pick-and-place operations [4].

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 [5],[6]. A single policy was proposed in [7] 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 [8] 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 [9], which is known to be NP-hard. A dynamic version of this problem was considered in [10], which also accounted for demands stochastically disappearing with a known distribution. The work in [11] 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 [12] and in [13]. 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 via a Poisson process in time with parameter λ. Each demand needs to be serviced exactly T time units after its generation. A demand is serviced by a vehicle if the vehicle is present at the demand location at its service time. 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.

The scenario of exact-time service with advance information arises in planning routes for a transportation service in which a demand for pick-up must be made T time units in advance. Exact time service requirements also arise in border patrol applications in which autonomous vehicles seek to intercept as many targets as possible that are moving towards a region [14]. An approach to this application is to allocate a subset of the border to each robot. Intruders arrive over time, attempting to cross this border. The time T then represents an estimate of the time between when the crossing point of an intruder becomes known, and when the attempted crossing will occur. The robots goal is to maximize the number of intruders that are successfully intercepted at their crossing point. Another application of this work is related to the scenario of catching the maximum number of balls thrown at a robot [15]. In this application, the robot has three degrees of freedom, but its reachable workspace is approximated by a planar, non-convex region. The work presented in this paper provides an initial approximation to these real-world problems, and we believe it provides some interesting insights into the structure of these problems.

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 T and μ_{max}, which is the maximum time taken by the vehicle to traverse the diameter of the environment, 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 has been studied in literature under the terminology of *competitive ratio*, e.g., see [16] 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 $\mu_{\text{max}}/T \to 0^+$. 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 $T \geq \mu_{\text{max}}$, we derive a second bound, which is an explicit function of μ_{max} and λ, and therefore, useful as a design tool. In particular, given a demand generation rate λ, this second bound can be used to determine the minimum velocity (or acceleration, depending upon the vehicle motion model) with which the vehicle should move in order to guarantee a specified service fraction. 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.

The primary insight provided in this paper is how advance temporal information on location of demands can be converted into an underlying reachability graph in a space-time environment. This directed graph happens to be acyclic and, therefore, longest paths inside it can be computed efficiently in polynomial time. A preliminary version of this work was addressed in [17], 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 compact environment $\mathcal{E} \subset \mathbb{R}^2$, that contains a single vehicle whose position is denoted by $\mathbf{p}(t) = [X(t), Y(t)]^T \in \mathcal{E}$. We assume that the motion model admits a function $\mu : \mathcal{E} \times \mathcal{E} \to \mathbb{R}_{\geq 0}$ which denotes the minimum time taken for the vehicle to move from rest at a point in $\mathcal{E}$ to reach another point in $\mathcal{E}$ at rest. By being at rest, we mean that all higher derivatives of the position, i.e., velocity, acceleration, etc., are equal to zero. Such a function exists for single integrator and double integrator motion models, which we will address in detail later.

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

We would like to emphasize that our solution approach presented in this paper is *independent* of the above assumption about the demand arrival process. The assumptions on the arrival process are required only for deriving the analytic guarantees on performance of our solution. In particular, the Poisson arrival in time assumption is commonly employed in operations research and queueing theory literatures, and the uniform distribution is reasonable for scenarios in which there is no prior information on where the demands are more likely to appear, e.g., see [5], [18].

Let $\mathcal{Q}(t) \subset \mathcal{E}$ denote the set of positions of all released but unserviced demands at time t. If the ith demand is either served or is missed, then it is removed from $\mathcal{Q}$. The ith demand is *active* if it is released and is neither serviced nor missed.

Online and Offline Algorithms: An online algorithm [19] (or policy) for the vehicle is a map $\mathcal{P} : \mathcal{E} \times \mathbb{F}(\mathcal{E}) \to \mathbb{R}^2$, 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 $t \mapsto p(t)$ 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 throughout the execution of our policy, 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

$$\mathbb{F}_{\text{cap}}(P) := \limsup_{t \to +\infty} \mathbb{E}\left[\tfrac{m_{\text{srv}}(t)}{m_{\text{srv}}(t) + m_{\text{miss}}(t)}\right], \tag{1}$$

where $m_{\text{srv}}(t), m_{\text{miss}}(t)$ are the number of demands that are served and missed until time t, respectively, and the expectation is with respect to the stochastic process that generates the demands.

B. Background Results

Since our goal is to characterize the service fraction, we will leverage existing results on the distribution of demands in an unserviced region to yield a bound on the average number of demands missed per single demand serviced. Then, we will review the longest path problem as we will use it to define our policies on an appropriately defined directed acyclic graph.

1) Demand distribution: Suppose that demands are generated uniformly randomly in a planar region $\mathcal{E} \in \mathbb{R}^2$, and as per a Poisson process in time with rate λ. Then, for every positive integer i, the ith demand is represented uniquely by a triple $(x_i, y_i, t_{\text{rel},i}) \in \mathcal{E} \times \mathbb{R}_{\geq 0}$. The following result characterizes the distribution of demands within any region contained in $\mathcal{E} \times [0, t]$.

Lemma II.1 (Distribution of outstanding demands, [20]) *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* $\mathcal{E} \times [0, t]$,

$$\mathbb{P}[|\mathcal{R} \cap \mathcal{Q}| = n] = \frac{\mathrm{e}^{-\lambda V / A(\mathcal{E})} (\lambda V / A(\mathcal{E}))^n}{n!}, \tag{2}$$

where $A(\mathcal{E})$ *is the area of* $\mathcal{E}$. *As a consequence, conditioned on the number of demands within such a region* $\mathcal{R}$, *the demands are distributed uniformly randomly in* $\mathcal{R}$.

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 [21]. However, if the graph is a Directed Acyclic Graph (DAG), then the longest path problem has an efficient dynamic programming solution [22] with complexity $O(|V| + |E|)$, that relies on topologically sorting [23] 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 ith demand is active at time t, the time for which the demand remains active is given by

$$t_{\text{rel},i} + T - t, \tag{3}$$

where $t_{\text{rel},i}$ is the generation time of the demand. 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 $t_{\text{rel},i} \leq t \leq t_{\text{rel},i} + T$. Recall that from the assumption on the vehicle dynamics, the function $\mu(\mathbf{p}(t), \mathbf{q}_i)$ denotes the minimum time taken for the vehicle to move from rest at $\mathbf{p}(t)$ to reach $\mathbf{q}_i$ at rest. Clearly, demand i can be serviced if and only if

$$\mu(\mathbf{p}(t), \mathbf{q}_i) \leq t_{\text{rel},i} + T - t. \tag{4}$$

From this description, we can define the set of reachable demands from any location in $\mathcal{E}_{\text{ST}}$.

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*

$$R_T(\mathbf{y}, t) := \{(\mathbf{z}, \tau) \in \mathcal{E} \times [t, T] \mid \mu(\mathbf{y}, \mathbf{z}) \leq \tau - t\}. \tag{5}$$

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 $(\mathbf{q}_i, t) = (x_i, y_i, t) \in R_T(\mathbf{p}(\bar{t}), \bar{t})$. 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. In the space-time representation, every demand that is reachable from $(\mathbf{q}_i, t_i)$, is reachable from $(\mathbf{p}(\bar{t}), \bar{t})$. 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 $\mathcal{E}_{\text{ST}}$. This idea motivates the concept of a reachability graph in $\mathcal{E}_{\text{ST}}$.

Definition III.2 (Reachability graph) *The reachability graph of a set of points $\{(\mathbf{q}_1, t_1), \ldots, (\mathbf{q}_n, t_n)\} \in \mathcal{E}_{\text{ST}}$, is a directed acyclic graph with vertex set $V := \{1, \ldots, n\}$, 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 $\mathcal{Q}$ of n outstanding demands, and a vehicle position $(\mathbf{p}, t)$, we can compute the corresponding reachability graph. For the purpose of illustration, we will represent the reachability graph assuming a simple first-order integrator motion with bounded speed for the vehicle (cf. Fig. 1). However, the concept generalizes to higher order dynamical models since the vehicle needs to start and end its motion from the rest condition while moving between demands.

In addition, by Section II-B2, we can compute the longest path in a reachability graph in $O(n^2)$ computation time.

Remark III.3 (Finite onsite service times) The reachability graph concept can be extended to the case of finite onsite service time equal to Δt as follows. For every demand i, create a copy of it at time instant $t_{\text{rel},i} + T + \Delta t$. Create a directed edge from the demand location to that of its copy. Now, demand j is reachable from i if and only if $(\mathbf{q}_j, t_{\text{rel},j}) \in R_T(\mathbf{q}_i, t_{\text{rel},i} + T + \Delta t)$. For the ease of presentation, we will consider zero onsite service time in this paper.

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. Notice that at least one demand is serviced on the path prior to recomputation of the longest path. This excludes the possibility of undesired switching behavior, where a vehicle continuously switches between paths without ever servicing any.

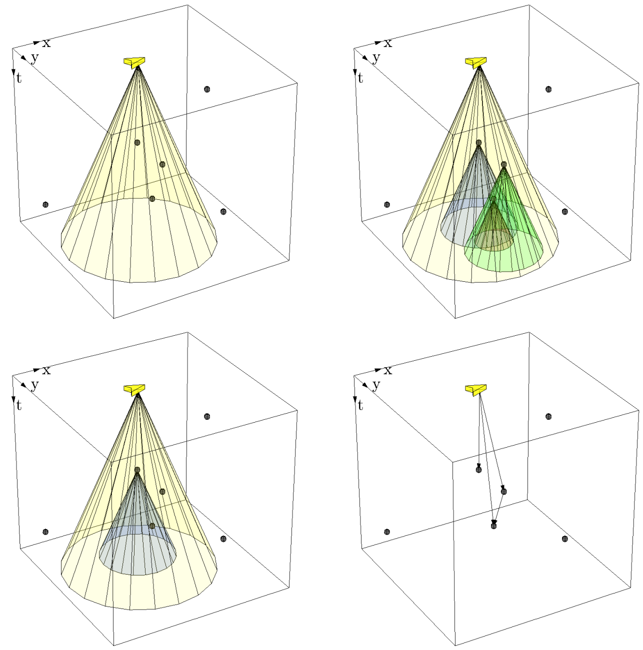

Fig. 1. The construction of the reachability graph. The top-left figure shows the set of reachable demand locations in the space-time environment $\mathcal{E}_{\text{ST}}$ from the vehicle positioned at a location in the environment $\mathcal{E}$. The top-right and bottom-left figures show the reachable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

Algorithm 1: The Longest Path (LP) policy

1 Compute the reachability graph of the vehicle position and all unserviced demands in $\mathcal{Q}(0)$.
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 the fraction of demands served on the path exceeds $\eta \in]0, 1]$, 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 literature, such a policy is referred to as an *offline algorithm* [19]. 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 $\mathcal{Q}(0) \cup \bar{\mathcal{Q}}$.
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.

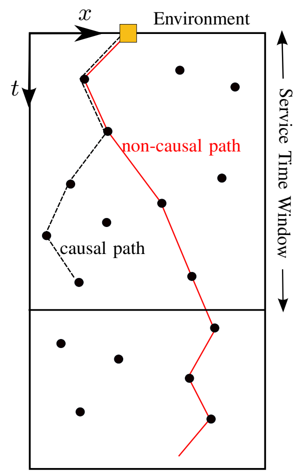

Fig. 2. A snapshot in the evolution of the Non-causal Longest Path policy as shown in the $\mathcal{E}_{\text{ST}}$ viewed along the $+y$ direction. The vehicle has planned the solid red path through all demands, including those that have not yet arrived. In comparison, a dashed causal longest path is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

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 $\mathcal{Q}(t)$. 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 [19]. Define,

$$\mu_{\max} := \max_{\mathbf{y},\mathbf{z}\in\mathcal{E}} \mu(\mathbf{y},\mathbf{z}), \tag{6}$$

which is the maximum of the time taken by the vehicle to move between any two points in the environment. Then, the following result holds.

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

$$\mathbb{F}_{\text{cap}}(\text{LP}) \geq \left(1 - \frac{\mu_{\max}}{T}\right)\mathbb{F}_{\text{cap}}(\text{NCLP}), \tag{7}$$

so that the Longest Path policy is optimal as $\mu_{\max}/T \to 0^+$.

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).

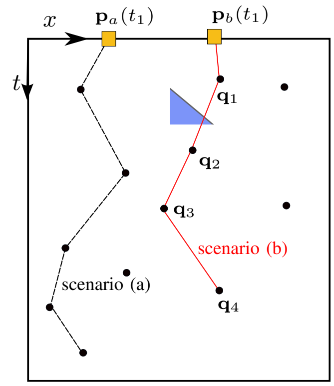

Fig. 3. Scenario (a) and (b) for the proof of Theorem IV.1, as viewed along the $+y$ direction. Path (a) visits five demands and thus $\mathcal{L}_a = 5$. Path (b) visits four demands, yielding $m = 4$. The demand $\mathbf{q}_2$ is the highest on path (b) that can be serviced from $\mathbf{p}_a(t_1)$.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

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 $(\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_m) \in \mathcal{Q}(t_1)$. The demand $\mathbf{q}_1$ is reachable from $\mathbf{p}_b(t_1)$, but it may not be reachable from $\mathbf{p}_a(t_1)$. To obtain a lower bound on the length of the longest path in scenario (a), consider the set of demands $(\mathbf{q}_{n+1}, \mathbf{q}_{n+2}, \ldots, \mathbf{q}_m)$, where $\mathbf{q}_{n+1}$ is the earliest demand on the path $(\mathbf{q}_1, \ldots, \mathbf{q}_m)$ that can be reached from $\mathbf{p}_a(t_1)$, see Fig. 3. Clearly this set has cardinality that can be at most the length of the longest path in scenario (a). Thus, the length of the longest path in scenario (a), $\mathcal{L}_a$, is at least

$$\mathcal{L}_a \geq m - n, \tag{8}$$

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

Now, the vehicle in scenario (a) can service any demand $(\mathbf{q}_i, t_{\text{rel},i})$ with $t_{\text{rel},i} + T - t_1 \geq \mu_{\max}$. Thus, the demands $\mathbf{q}_1, \ldots, \mathbf{q}_n$ must all have their release times in $[t_1 - T, t_1 + \mu_{\max} - T]$, since they are active at time t_1. Let the total number of active demands at time t_1 be N_{tot}. Then, conditioned on N_{tot}, by Lemma II.1, the active demands are distributed uniformly randomly in the region $\mathcal{E} \times [t_1 - T, t_1]$. Conditioned on N_{tot}, the density of the active demands in this region is $N_{\text{tot}}/(A(\mathcal{E})T)$, where $A(\mathcal{E})$ is the area of the environment. Since $\mu_{\max} \leq T$, the expected number of active demands contained in the region $\mathcal{E} \times [t_1 - T, t_1 + \mu_{\max} - T]$ is $N_{\text{tot}} A(\mathcal{E})\mu_{\max}/(A(\mathcal{E})T) = N_{\text{tot}}\mu_{\max}/T$. Hence,

$$\mathbb{E}\left[n|N_{\text{tot}}\right] = N_{\text{tot}}\frac{\mu_{\max}}{T}\mathbb{F}_{\text{cap}}(\text{NCLP}). \tag{9}$$

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

$$\mathbb{E}\left[m|N_{\text{tot}}\right] = N_{\text{tot}}\mathbb{F}_{\text{cap}}(\text{NCLP}). \tag{10}$$

Combining equations (9) and (10) with equation (8) we obtain

$$\mathbb{E}\left[\mathcal{L}_a|N_{\text{tot}}\right] \geq N_{\text{tot}}\left(1 - \frac{\mu_{\max}}{T}\right)\mathbb{F}_{\text{cap}}(\text{NCLP}), \tag{11}$$

$$\mathbb{E}\left[\tfrac{\mathcal{L}_a}{N_{\text{tot}}}|N_{\text{tot}}\right] \geq \left(1 - \frac{\mu_{\max}}{T}\right)\mathbb{F}_{\text{cap}}(\text{NCLP}). \tag{12}$$

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

$$\mathbb{E}\left[\tfrac{\mathcal{L}_a}{N_{\text{tot}}}\right] = \mathbb{E}\left[\mathbb{E}\left[\tfrac{\mathcal{L}_a}{N_{\text{tot}}}|N_{\text{tot}}\right]\right] \geq \left(1 - \frac{\mu_{\max}}{T}\right)\mathbb{F}_{\text{cap}}(\text{NCLP}). \tag{13}$$

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 $\mathbb{F}_{\text{cap}}(\text{LP}) \geq \mathbb{E}\left[\mathcal{L}_a/N_{\text{tot}}\right]$ 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 Path policy, presented in Algorithm 3.

Algorithm 3: The Greedy Path (GP) policy

Assumes: Vehicle is located at $(X(t), Y(t))$

Compute the reachability set $R_T((X(t), Y(t)), t)$.
Service the demand in $R_T((X(t), Y(t)), t)$ with the highest time coordinate.
If no demands exist, move toward the center of the environment.
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 $\mathbf{p}_a(t), \mathbf{p}_b(t) \in \mathcal{E}$ denote the positions of the vehicle when following the Longest path and the Greedy path policies from the initial instant, respectively. The capture fraction $\mathbb{F}_{\text{cap}}$ can be viewed as a function of $\mathcal{Q}(t)$ and the vehicle position along with the choice of policy P. Since both $\mathbf{p}_a(t)$ and $\mathbf{p}_b(t)$ are independent of $\mathcal{Q}(t)$,

$$\mathbb{F}_{\text{cap}}(LP, \mathcal{Q}(t), \mathbf{p}_a(t)) = \mathbb{F}_{\text{cap}}(LP, \mathcal{Q}(t), \mathbf{p}_b(t)). \tag{14}$$

But conditioned on the position $\mathbf{p}_b(t)$,

$$\mathbb{F}_{\text{cap}}(LP, \mathcal{Q}(t), \mathbf{p}_b(t)) \geq \mathbb{F}_{\text{cap}}(GP, \mathcal{Q}(t), \mathbf{p}_b(t)), \tag{15}$$

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

$$\mathbb{F}_{\text{cap}}(LP, \mathcal{Q}(t), \mathbf{p}_a(t)) \geq \mathbb{F}_{\text{cap}}(GP, \mathcal{Q}(t), \mathbf{p}_b(t)), \tag{16}$$

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.2 (Lower Bound for Longest Path policy) *If $T \geq \mu_{\max}$, then for the Longest Path policy*

$$\mathbb{F}_{\text{cap}}(\text{LP}) \geq \mathbb{F}_{\text{cap}}(\text{GP}) \geq 1/g(\lambda, \mu_{\max}, \mathcal{E}), \tag{17}$$

where the function

$$\begin{aligned} g(\lambda, \mu_{\max}, \mathcal{E}) := 1 + \lambda\Big[&\int_0^{\mu_{\max}} \left(t - \int_0^t \rho(Z_\tau) d\tau\right) f(t) dt \\ &+ \left(\mu_{\max} - \int_0^{\mu_{\max}} \rho(Z_t) dt\right) \mathrm{e}^{-\lambda \int_0^{\mu_{\max}} \rho(Z_t) dt}\Big], \end{aligned} \tag{18}$$

the set $Z_t \subseteq \mathcal{E}$ denotes the set of points reachable from the worst-case service vehicle location in $\mathcal{E}$ in time $t \in [0, \mu_{\max}]$, $A(\cdot)$ denotes the area of a planar region, the scalar $\rho(Z_t) := A(Z_t)/A(\mathcal{E})$ is the ratio of the area of the reachable set Z_t to that of the environment, and the function

$$f(t) = \lambda \rho(Z_t) \mathrm{e}^{-\lambda \int_0^t \rho(Z_\tau) d\tau}. \tag{19}$$

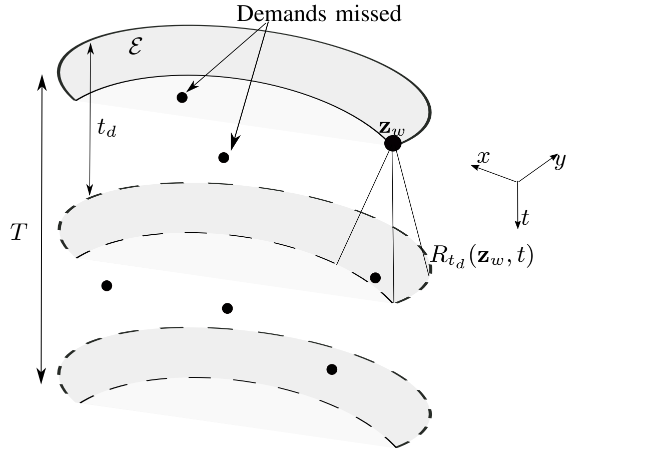

Fig. 4. The setup for the proof of Theorem IV.2. The service vehicle is located at $(\mathbf{z}_w, \bar{t})$. The circle centered at $\mathbf{z}_w$ with radius equal to t_d is the set of points which the vehicle can reach at rest starting from rest at $\mathbf{z}_w$. The worst case for the region $R_t(\mathbf{z}_w, \bar{t})$ is indicated by the green solid. As per the Greedy policy, any demand within $\mathcal{E}_{\text{ST}}$, which is outside of the green solid region is left unserviced.

Proof: Notice that if $m_{\text{srv}}(t) > 0$ for some $t > 0$, then

$$\begin{aligned} \limsup_{t \to +\infty} \mathbb{E}\left[\frac{m_{\text{srv}}(t)}{m_{\text{srv}}(t) + m_{\text{miss}}(t)}\right] &= \limsup_{t \to +\infty} \mathbb{E}\left[\frac{1}{1 + \frac{m_{\text{miss}}(t)}{m_{\text{srv}}(t)}}\right] \\ &\geq \left(1 + \limsup_{t \to +\infty} \mathbb{E}\left[\frac{m_{\text{miss}}(t)}{m_{\text{srv}}(t)}\right]\right)^{-1}, \end{aligned} \tag{20}$$

where the last step comes from an application of Jensen's inequality [24]. 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 $\bar{t}$ at which the service vehicle services its ith 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 $\mathcal{E}_{\text{ST}}$; namely, the position $\mathbf{z}_w \in \mathcal{E}$ from which the reachability set $R_T(\mathbf{z}_w, t)$ has the least volume, for every $T > 0$. An illustration of this location for a crescent-like, non-convex environment (also considered later in Section V), is shown in Figure 4.

Let $Z_t \subset \mathcal{E}$ denote the set of points that can be reached from $\mathbf{z}_w$ in time $t \in [0, \mu_{\max}]$. This set is the intersection of the purple region with $\mathcal{E}$ in Fig. 4. Let $A(Z_t)$ denote the area of this set and let $|R_t|$ denote the volume of $R_t(\mathbf{z}_w, \bar{t})$. An illustration of the set $R_t(\mathbf{z}_w, \bar{t})$ is shown as a green solid in Fig. 4. Then,

$$|R_t(\mathbf{z}_w, \bar{t})| = \begin{cases} \int_0^t A(Z_\tau) d\tau, & \text{if } z \leq \mu_{\max}, \\ \int_0^{\mu_{\max}} A(Z_\tau) d\tau + A(\mathcal{E})(t - \mu_{\max}), & \text{if } t > \mu_{\max}. \end{cases} \tag{21}$$

Let t_d be the time for which the first reachable demand is active at time $\bar{t}$. That is,

$$t_d := \min_{(x,y,t) \in \mathcal{Q}(\bar{t}) \cap R_t(\mathbf{z}_w, \bar{t})} \{T - t\}, \tag{22}$$

where $\mathcal{Q}(\bar{t})$ is the set of outstanding demands at time $\bar{t}$. By Lemma II.1, the probability that a subset $\mathcal{B} \subset \mathcal{E}$ with volume $|\mathcal{B}|$ contains zero demands is given by

$$\mathbb{P}[|\mathcal{B} \cap \mathcal{Q}(\bar{t})| = 0] = \mathrm{e}^{-\lambda |\mathcal{B}|/A(\mathcal{E})}, \tag{23}$$

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

$$\mathbb{P}[t_d > t] = \mathbb{P}[|R_t(\mathbf{z}_w, \bar{t}) \cap \mathcal{Q}(\bar{t})| = 0] = \mathrm{e}^{-\lambda |R_t(\mathbf{z}_w, \bar{t})|/A(\mathcal{E})}. \quad (24)$$

The probability density function of t_d for $t \leq \mu_{\max}$, which is obtained from the expression

$$f(t) = \frac{d}{dt}(1 - \mathbb{P}[t_d > t]) = -\frac{d}{dt}\mathrm{e}^{-\lambda |R_t(\mathbf{z}_w, \bar{t})|/A(\mathcal{E})}, \quad (25)$$

satisfies

$$f(t) = \frac{\lambda A(Z_t)}{A(\mathcal{E})}\mathrm{e}^{-\lambda \int_0^t A(Z_\tau)d\tau/A(\mathcal{E})}. \quad (26)$$

Now, given t_d, all demands residing in the region $\text{miss}_{t_d} := (\mathcal{E} \times [T - t_d, T]) \setminus R_t(\mathbf{z}_w, \bar{t})$ will be lost unserviced (see Fig. 4). Using the expressions for $|R_t(\mathbf{z}_w, \tau)|$, the volume of miss_{t_d} is given by

$$|\text{miss}_{t_d}| = \begin{cases} t_d A(\mathcal{E}) - \int_0^{t_d} A(Z_\tau)d\tau, & \text{if } t_d \leq \mu_{\max}, \\ A(\mathcal{E})\mu_{\max} - \int_0^{\mu_{\max}} A(Z_\tau)d\tau, & \text{if } t_d > \mu_{\max}. \end{cases} \quad (27)$$

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

$$\begin{aligned} \mathbb{E}\left[m_{\text{miss},i}\right] &= \frac{\lambda}{A(\mathcal{E})}\mathbb{E}\left[|\text{miss}_{t_d}|\right] \\ &= \frac{\lambda}{A(\mathcal{E})}\Big[\int_0^{\mu_{\max}} \left(tA(\mathcal{E}) - \int_0^t A(Z_\tau)d\tau\right) f(t)dt \\ &\quad + \left(A(\mathcal{E})\mu_{\max} - \int_0^{\mu_{\max}} A(Z_t)dt\right)\mathbb{P}[t_d > \mu_{\max}]\Big]. \quad (28) \end{aligned}$$

Since $\mathbb{E}\left[m_{\text{miss},i}\right]$ is computed for the worst-case vehicle position $(\mathbf{z}_w, \bar{t})$ in $\mathcal{E}_{\text{ST}}$, and since this expression holds at every demand service, by equation (20) and using equation (24), we obtain the desired result. ■

Remark IV.3 (Asymptotic performance) Our bounds achieve optimality in the asymptotic regimes of $\mu_{\max}/T \to 0^+$ for the first bound, and $\lambda \to 0^+$ or $\mu_{\max} \to 0^+$ for the second bound.

C. Special Case: Single Integrator Dynamics & Square Environment

In this section, we show how the analytic results reduce to closed-form expression under special circumstances. Specifically, we consider *first-order integrator dynamics with maximum speed* u, and *square operating environments with edge length* W.

Corollary IV.4 (Simple motion and square environments) *For simple vehicle motion with maximum speed u and in square environments with edge length W,*

(i) *the capture fraction of the Longest path and the non-causal Longest Path satisfy*

$$\mathbb{F}_{\text{cap}}(\text{LP}) \geq \left(1 - \frac{\sqrt{2}W}{Tu}\right)\mathbb{F}_{\text{cap}}(\text{NCLP}). \quad (29)$$

(ii) *If $T \geq \sqrt{2}W/u$, then the Longest Path policy satisfies*

$$\mathbb{F}_{\text{cap}}(\text{LP}) \geq \mathbb{F}_{\text{cap}}(\text{GP}) \geq 1/g(\lambda, u, W), \quad (30)$$

with the function g defined as

$$g := \mathrm{e}^{-\frac{\sqrt{2}\lambda W}{3u}} + \frac{\lambda}{3}\Big(\frac{6W^2}{\lambda u^2}\Big)^{\frac{1}{3}}\Big(\Gamma\big(\frac{1}{3}, 0\big) - \Gamma\big(\frac{1}{3}, \frac{\sqrt{2}\lambda W}{3u}\big)\Big), \quad (31)$$

and the function

$$\Gamma(a, z) := \int_z^{+\infty} t^{a-1}\mathrm{e}^{-t}dt, \quad (32)$$

is the incomplete Gamma function*.*

The first claim follows by direct substitution of $\mu_{\max} = \sqrt{2}W/u$. The second claim follows from the fact that $A(Z_z)$ has a closed form expression in terms of u and W in square environments, and the resulting integral can be evaluated in closed-form.

V. Simulations

We now present results from numerical experiments. Through these experiments, we compare the Longest Path policy, with the parameter $\eta = 1$, with the Non-causal Longest Path policy and to the theoretical lower bound in Theorem IV.2, in order to verify the analytic claims.

A. Simple motion and square environment

In the first experiment, we considered a square environment with edge length equal to W. The vehicle motion model was chosen to be simple first-order integrator with bounded speed u. To simulate the LP and the NCLP policies, we performed 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. We also confirm the analytic results, i.e., the bound relative to the non-causal Longest path policy and the explicit lower bound summarized in Corollary IV.4.

In this experiment, we also simulated the Longest Path policy for a variation of the problem with a service time window instead of exact service time. In other words, suppose that for every i, the ith demand could be serviced at any time within the time interval $[t_{\text{rel},i}, t_{\text{rel},i}+T]$. Then, the vehicle may still use the Longest Path policy to service the demands. The solid blue curve in Fig. 5 shows that the empirically observed performance is very close to that of the Longest Path policy in the exact time service scenario. By applying the Longest Path policy to the problem with service time windows, the analytic lower bounds (Theorem IV.1 and Corollary IV.4) still apply, as seen from Fig. 5. However, for large values of T, the longest path policy would prove to be very sub-optimal, as it corresponds to the vehicle simply waiting at a demand's location before servicing it. In this parameter regime, the performance may be improved using other heuristics that are based on the order in which the vehicle serves existing demands. The main challenge here is that history-based searches have much higher computational complexity than the longest path policy. We believe such algorithmic questions are a promising direction for future research.

B. Double-integrator motion in a non-convex environment

We now report the results of simulating the LP and the NCLP policies in a more complex, non-convex environment similar to the crescent-like, approximation of the reachable workspace of the robot in [15]. We consider the following scenario, illustrated in Figure 6. If $B(x, y, R)$ denotes the closed circular region centered around the point (x, y) in the plane. Then, our workspace is the closure of the region $B(0,0,R) \setminus (B(0,0,R) \cap B(-d,0,R))$, for some $d \in [0, R]$. In our experiments, we chose $R = 1$ and $d = R/2$. The vehicle motion model was chosen to be a double integrator with the control input u being the acceleration of the vehicle. This can be extended to a more detailed motion model for the vehicle that enables at least an approximate computation of the bound in Theorem IV.2. Balls are assumed to be moving with identical speeds v in a direction

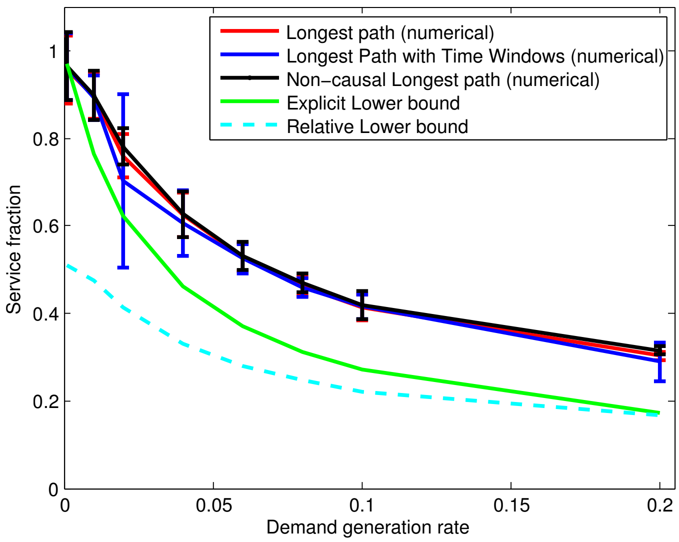

Fig. 5. Simulation results for LP policy (solid red line with error bars showing $\pm$ one standard deviation) and the NCLP policy (dashed black line) for a square environment of width $W = 100$ and maximum vehicle speed $u = 3$. This yields $T = 100 > \sqrt{2}W/u$, and the explicit lower bound in Theorem IV.2 is shown in solid green, while the lower bound relative to the non-causal Longest path is shown in solid cyan. The solid blue curve shows the performance of the Longest Path policy (recomputed after every serviced demand) for the scenario of finite service time window.

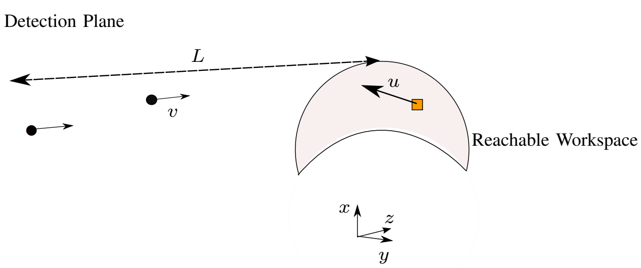

Fig. 6. An abstraction of the ball-catching scenario from [15]. The yellow square denotes the vehicle's location, the environment is crescent-like, non-convex planar region. Balls (denoted by black dots) are assumed to be moving with identical speeds v in a direction perpendicular to the environment. A ball is detected after it passes the shaded detection plane located at a distance L from the environment along the $-z$ direction.

perpendicular to the environment. A ball is detected after it passes the detection plane located at a distance L from the environment along the $-z$ direction. This problem can now be modeled in our current framework with the choice of $T = L/v$.

To simulate the LP and the NCLP policies, we performed 10 runs of the policy, where each run consists of 100 demands. A comparison of the service fractions for the two policies for a fixed demand generation rate λ and at varying vehicle acceleration values u is presented in Fig. 7. When $\mu_{\max} < T$, we observe that the service fraction of the LP policy is nearly identical to that of the NCLP policy, as shown in Fig. 7. Further, 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.2, are fairly consistent with the experimental results.

These results (both numerical and analytic) can be used to design the vehicle acceleration u, given a certain desired value of the service fraction. For example, for a generation rate of $\lambda = 0.5$, if the desired service fraction is 0.6, then the analytic result states that a vehicle

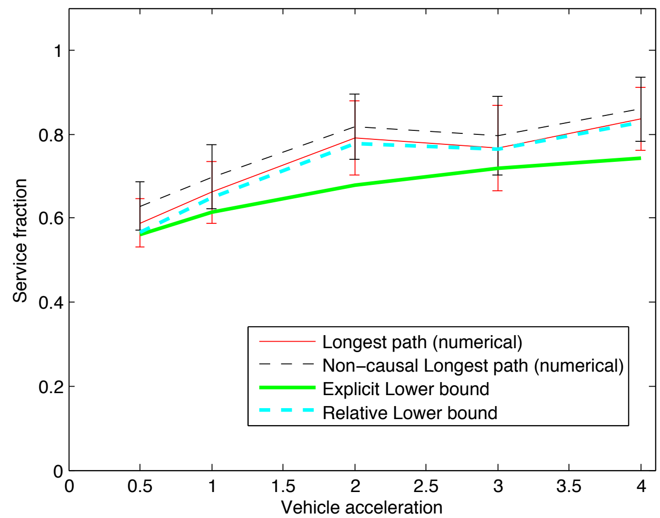

Fig. 7. Simulation results for LP policy (solid red line with error bars showing $\pm$ one standard deviation) and the NCLP policy (dashed black line) for a non-convex environment similar to the workspace of the robot in [15] and for a fixed demand generation rate $\lambda = 0.5$. A double integrator model is assumed in this experiment. The explicit lower bound, computed approximately using the expression in Theorem IV.2, is shown in solid green, while the lower bound relative to the non-causal Longest path from Theorem IV.1 is shown in dashed cyan.

acceleration of $u \geq 0.75$ is sufficient to achieve the goal, while the numerical result suggests that $u \geq 1$ is sufficient for the same. This gap is higher for higher values of the desired service fraction.

C. Key steps towards implementation

We believe that the work in this paper provides a good initial approximate solution for a realistic scenario, such as motion planning for the ball-catching robot. We identify the following gaps that need to be addressed before our work can be applied to real experimental scenarios:

(i) *Control action computation*: Our results are presented in terms of the function μ, which characterizes the time it takes to move the robot from one location to another. A closed form expression for such a function is not readily available for more complex robots, beyond the class of linear systems. Further, the computation of control actions, e.g., the precise actuator signals at the joints which would realize the motion of the robot from one point to the other, would need to be obtained through heuristics such as sampling based approaches. This aspect is a focus of future research.

(ii) *Effect of uncertainty in the time T*: Our present work neglects the effect of uncertainty in the time T at which each demand needs to be served. Uncertainty in real scenarios can arise due to several reasons ranging from obtaining precise information about demand's location to computation of the longest path approach coupled with the control action for the robot actuators. However, we believe that our approach based on the longest path can yield a good initial solution which may be further optimized for the more-realistic model of the overall system.

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 primary insight provided in this paper is how advance temporal information on location of demands can be converted into an underlying reachability graph in a space-time environment. This graph happens to be directed and acyclic, and therefore, we can compute longest paths in an efficient (polynomial time) manner. A novel policy based on repeated computation of longest paths through the available set of demands was proposed. 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 motion models in which the vehicle cannot be brought to rest at a demand location, e.g., Dubins model, and multi-vehicle versions of this problem.

REFERENCES

[1] R. Szechtman, M. Kress, K. Lin, and D. Cfir, "Models of sensor operations for border surveillance," *Naval Research Logistics*, vol. 55, no. 1, pp. 27–41, 2008.

[2] N. Basilico and S. Carpin, "Online patrolling using hierarchical spatial representations," in *IEEE Int. Conf. on Robotics and Automation*, St. Paul, MN, USA, May 2012, pp. 2163–2169.

[3] R. N. Smith, M. Schwager, S. L. Smith, B. H. Jones, D. Rus, and G. S. Sukhatme, "Persistent ocean monitoring with underwater gliders: Adapting sampling resolution," *Journal of Field Robotics*, vol. 28, no. 5, pp. 714–741, 2011.

[4] P. Chalasani and R. Motwani, "Approximating capacitated routing and delivery problems," *SIAM Journal on Computing*, vol. 28, no. 6, pp. 2133–2149, 1999.

[5] 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.

[6] ——, "Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles," *Operations Research*, vol. 41, no. 1, pp. 60–76, 1993.

[7] 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.

[8] 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.

[9] P. Toth and D. Vigo, Eds., *The Vehicle Routing Problem*, ser. Monographs on Discrete Mathematics and Applications. SIAM, 2001.

[10] 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.

[11] N. Bisnik, A. A. Abouzeid, and V. Isler, "Stochastic event capture using mobile sensors subject to a quality metric," *IEEE Transactions on Robotics*, vol. 23, no. 4, pp. 676–692, Aug. 2007.

[12] 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.

[13] J. Johnson and K. Hauser, "Optimal longitudinal control planning with moving obstacles," in *IEEE Intelligent Vehicles Symposium*, Jun. 2013, pp. 605–611.

[14] E. Raboin, P. Svec, D. Nau, and S. K. Gupta, "Model-predictive target defense by team of unmanned surface vehicles operating in uncertain environments," in *IEEE Int. Conf. on Robotics and Automation*, Karlsruhe, Germany, May 2013.

[15] J. Malzahn, A. S. Phung, and T. Bertram, "A multi-link-flexible robot arm catching thrown balls," in *7th German Conference on Robotics (ROBOTIK)*, Munich, Germany, May 2012.

[16] E. Koutsoupias and C. H. Papadimitriou, "On the k-server conjecture," *Journal of the ACM*, vol. 42, no. 5, pp. 971–983, 1995.

[17] 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.

[18] M. Pavone and E. Frazzoli, "Dynamic vehicle routing with stochastic time constraints," in *IEEE Int. Conf. on Robotics and Automation*, Anchorage, AK, May 2010, pp. 1460–1467.

[19] A. Borodin and R. El-Yaniv, *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.

[20] 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.

[21] B. Korte and J. Vygen, *Combinatorial Optimization: Theory and Algorithms*, 4th ed., ser. Algorithmics and Combinatorics. Springer, 2007, vol. 21.

[22] N. Christofides, *Graph Theory: An Algorithmic Approach*. Academic Press, 1975.

[23] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. MIT Press, 2001.

[24] L. Breiman, *Probability*, ser. Classics in Applied Mathematics. SIAM, 1992, vol. 7, corrected reprint of the 1968 original.