Dynamic Vehicle Routing in Presence of Random Recalls

Shaunak D. Bopardikar Vaibhav Srivastava

***Abstract*—Dynamic Vehicle Routing (DVR) problems involve a vehicle that seeks to service demands which are generated via a spatio-temporal stochastic process in a given environment. This paper introduces a DVR problem in which the vehicle needs to return to a central facility from time to time. We model the return events as a Poisson process with a known parameter. The problem parameters are the demand generation rate, the size of the environment and the recall rate. The goal is to design service policies for the vehicle in order to minimize the expected service time per demand. The contributions are as follows. We first provide a complete analysis of the regime of low demand arrival using a first-come-first-served policy. For the regime of high demand arrival, we derive a policy independent lower bound on the expected service time as a function of the problem parameters. We then adapt a well-known policy based on repeated computation of the Euclidean Traveling Salesperson tour through unserviced demands and provide an upper bound on the expected service time, quantifying the factor of optimality relative to the lower bound. We supplement the analysis with several insightful numerical simulations.**

***Index Terms*—Motion planning, stochastic processes, optimization**

I. Introduction

DYNAMIC vehicle routing (DVR) refers to a class of problems in which one or many vehicles seek to service demands that appear sequentially in a given environment. The goal is to determine in what sequence should each vehicle service demands in order to minimize the average time between the generation of a demand to the demand getting serviced. These problems arise in several applications ranging from surveillance and environmental monitoring [1] to efficient package delivery [2] using drones. This paper introduces a new variant of DVR problems in which a vehicle is required to periodically visit a given facility in the environment as per the outcome of a random process. This could be due to a high priority task to be completed at the facility or for re-fueling purposes when working in uncertain environments or for pro-active maintenance reasons.

Classic vehicle routing problems are concerned with planning optimal vehicle routes to visit a set of demands. The routes are planned with complete information of the targets and thus, the optimization is static but combinatorial [3]. In contrast, DVR considers scenarios in which the demand information is not known a priori, and thus routes must be re-planned as new information becomes available over time. This problem was introduced on graph environments in [4]. Fundamental limits, novel policies and their constant factor optimality guarantees in continuous environments were established in [5]. Subsequent developments in this line of work have been mostly for specific novel features in the problem formulation. Demands may be static, but in a dynamically varying environment [6], [7] or with multiple levels of importance [8]. The vehicle may be tasked with performing pickup and delivery operations [9], [10] or may possess motion constraints [11], [12], [13]. Extensions to multiple vehicles case which do not require explicit communication have also been considered [14]. The demands may be mobile as considered in our body of work [15], [16], [17]. We refer the reader to [18] for detailed review of this field. The present paper considers a DVR problem for a single vehicle with simple motion and demand arrival process similar to [19], but with a novel aspect of requiring the vehicle to visit a centrally located facility as per a temporal stochastic process. Key differences with [8] lie in the problem formulation wherein their goal was to minimize a *strict* convex combination of the waiting times of the high and low priority demands and therefore, the analysis tools from [8] are not applicable in the present setting.

Modeling the recall process can be an elaborate task that depends on the particular application, such as recall due to failure or for re-fueling purposes. For analytic tractability, this paper models recall as a random process, typically Poisson distributed with parameter equal to $\mu \geq 0$. The demands are generated uniformly randomly in the environment and as per another, independent Poisson process over time with a known parameter $\lambda > 0$. The problem parameters are the demand generation rate λ, the size of the environment modeled as a compact region of area A and the parameter μ of the recall model. The goal is to design service policies for the vehicle in order to minimize the expected time taken to service a demand.

Our contributions are as follows. We first analyze the parameter regime called *light load* ($\lambda \to 0^+$) using a policy based on first-come-first-served order of servicing the demands. We provide a closed form expression for the expected service time per demand in this regime. We then focus the majority of this paper analyzing the *heavy load* regime ($\lambda \to +\infty$). In this regime, we first derive a fundamental limit to the problem in terms of a lower bound on the expected service time as a function of the problem parameters which holds for *any admissible policy* for the vehicle. We then design a service policy by adapting a well-known policy based on repeated computation of the Euclidean Traveling Salesperson (ETSP) tour through unserviced demands. We provide an upper bound on the expected service time for this policy and show that it performs within a constant factor of the lower bound, where the constant factor depends only on A and μ. Finally, we numerically study the: 1) performance of the TSP-based policy in non-asymptotic parameter regimes; 2) comparison with a simple *nearest neighbor* policy and, 3) performance of the TSP-based policy with a recall model based on a linear hazard rate that generalizes the Poisson recall process considered in

This work was supported in part by NSF under Award IIS-1734272.

The authors are with the Electrical and Computer Engineering Department, Michigan State University, East Lansing, MI. Emails: {shaunak, vaibhav}@egr.msu.edu

the analysis.

This paper is organized as follows. Section II presents the mathematical description of the problem considered. Sections III and IV present service policies for the vehicle and their analysis for the light and the heavy load regimes, respectively. Section V presents a numerical study of the heavy load problem setting. Section VI summarizes this work and identifies directions for future extensions.

II. Problem Formulation and Background

We begin with the problem statement and present some useful background results.

A. Modeling and Problem Statement

We consider an environment $\mathcal{E} \subset \mathbb{R}^2$, assumed to be a compact convex region with area A. While it is possible to extend the analysis in this paper to higher dimensions, for ease of exposition, we focus on a 2 dimensional version in this paper wherein $\mathcal{E}$ is a square. The demands for service are generated according to a temporal Poisson process with parameter $\lambda > 0$, and their locations are independent and uniformly distributed in $\mathcal{E}$. Each demand requires an independent and identically distributed (i.i.d.) amount of on-site service time with mean $\bar{s} > 0$ and finite second moment. The *service time* of demand i, denoted as T_i is defined as the time elapsed between the arrival of the i-th demand and the time the vehicle services the i-th demand. The steady-state service time is $\mathbb{E}[T] := \limsup_{i\to+\infty} \mathbb{E}[T_i]$, which refers to the expected service time of a demand after a transient due to the initial location of the vehicle and the demand generation process has passed.

The service vehicle is modeled as a single integrator with speed normalized to unity. Thus, we will use time elapsed and distance traveled interchangeably. Specifically, let $\mathbf{p}(t) \in \mathcal{E}$ denote the position of the vehicle at time t. Let $\mathcal{Q}(t) \subset \mathcal{E}$ denote the set of all unserviced demand locations at time t, and $n(t)$ the cardinality of $\mathcal{Q}(t)$. Servicing of a demand $\mathbf{q}_i \in \mathcal{Q}$ and removing it from the set $\mathcal{Q}$ occurs when the vehicle reaches the location of the demand and spends an i.i.d. amount of on-site service time. A static feedback control policy for the system 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 vehicle as a function of the current state of the system, given by

$$\dot{\mathbf{p}}(t) = \mathcal{P}(\mathbf{p}(t), \mathcal{Q}(t)).$$

A policy $\mathcal{P}$ is said to be *stable* if the expected number of unserviced demands in the environment is uniformly bounded at all times. A standard calculation (cf. [19]) yields that if $\mathbb{E}[T]$ is finite for a policy, then the expected number of unserviced demands, $\mathbb{E}[n] = \lambda(\mathbb{E}[T] - \bar{s})$. The *load factor* $\rho := \lambda\bar{s}$, signifies the fraction of time the vehicle performs on-site service. A necessary condition for the existence of a stable policy is that the load factor $\rho < 1$, i.e., in the expected amount of time during which the vehicle performs on-site service, no new demand is expected to be generated.

The novel aspect introduced in this work is that the service vehicle has to return to a facility, located at the center of $\mathcal{E}$ periodically over time. The return events are assumed to be generated independently of the demand generation process and occur as per a Poisson process with intensity $\mu > 0$. *The goal is to design a policy $\mathcal{P}$ for the service vehicle to minimize $\mathbb{E}[T]$ in presence of the stochastic return process.* The expectation $\mathbb{E}[\cdot]$ is with respect to the joint distribution of the demand generation process and the recall process which are assumed to be independent.

Solution outline: We design service policies and present their analysis in both the light load regime ($\rho \to 0^+$ such that $\lambda \to 0^+$) and the heavy load regime ($\rho \to 1^-$ such that $\lambda \to +\infty$ and $\bar{s} \to 1/\lambda$). We will begin with the light load analysis and then present the heavy load results. In the heavy load, minimizing $\mathbb{E}[T]$ directly is difficult and so, we will first derive a lower bound on $\mathbb{E}[T]$ for any stable policy. We will then propose and analyze a policy which requires the vehicle to repeatedly compute the ETSP tour through the number of unserviced demands in the environment and then compare an upper bound on $\mathbb{E}[T]$ for the policy with the lower bound. For this purpose, in the next subsection, we will review some properties on length of the ETSP and inter-demand distance in presence of a large number of demands.

B. Shortest path through a large number of points

Given a set $\mathcal{Q}$ of n points in $\mathbb{R}^2$, the ETSP problem is to determine the shortest tour, i.e., a closed path that visits each point exactly once. Let $\mathrm{ETSP}(\mathcal{Q})$ denote the length of the ETSP tour through $\mathcal{Q}$. Below is classic result from [20].

Theorem II.1 (Asymptotic ETSP length, [20]) *If a set $\mathcal{Q}$ of n points are distributed independently and uniformly in a compact region of area A, then there exists a constant β_{TSP} such that, almost surely,*

$$\lim_{n\to+\infty} \frac{\mathrm{ETSP}(\mathcal{Q})}{\sqrt{n}} = \beta_{\mathrm{TSP}}\sqrt{A}. \tag{1}$$

The constant β_{TSP} has been estimated numerically as $\beta_{\mathrm{TSP}} \approx 0.7120 \pm 0.0002$.

The following is a useful intermediate result providing a lower bound on the expected distance to travel between demands [21, Page 23].

Theorem II.2 (Travel distance lower bound, [21]) *In the limit of $\rho \to 1^-$, if the expected number of unserviced demands in $\mathcal{E}$ is $\mathbb{E}[n]$, then a lower bound on the expected travel distance between demands is*

$$\frac{\beta_{\mathrm{TSP}}}{\sqrt{2}}\sqrt{\frac{A}{\mathbb{E}[n]}}.$$

III. Light Load Analysis

This section formalizes a policy for the vehicle in the light load, i.e., $\rho \to 0^+$ such that $\lambda \to 0^+$. The idea is that the vehicle remains at the facility location until there is a demand in the environment. While there are unserviced demands, the vehicle serves them in a *first-come-first-served (FCFS) manner*. If a recall event is generated or if there are

no unserviced demands, the vehicle returns to the facility and then resumes the FCFS order.

The following is a simple result that quantifies the expected time taken to service a demand under the stochastic return model considered in this paper.

Lemma III.1 (Travel from facility to a demand) *Suppose that the vehicle is located at the facility. Under the stochastic return model considered, the expected distance D_0 traveled by the vehicle to any demand which is at a distance of c from the facility is*

$$D_0 = \frac{2}{\mu}(\mathbf{e}^{\mu c} - 1) - c,$$

where the expectation is with respect to the return process.

Proof: Without any loss of generality, assume that the demand is along the X axis from the facility and the vehicle is at a distance $x \in \mathbb{R}_{\geq 0}$ from the facility along the shortest path from the facility to the demand, which is at a distance $c \in \mathbb{R}_{\geq 0}$ from the facility. Let $D(x;c) \mapsto \mathbb{R}_{\geq 0}$ be the expected distance the vehicle travels starting from a location at distance x from the facility to the demand at distance c. Note that $D_0 = D(0;c)$. Over a small incremental displacement Δx towards the demand, the following recursion holds:

$$\begin{aligned} D(x;c) &= \mu\Delta x \exp(-\mu\Delta x)(D_0 + x) \\ &\quad + \exp(-\mu\Delta x)(\Delta x + D(x+\Delta x;c)) + o(\Delta x), \end{aligned} \quad (2)$$

where $\mu\Delta x \exp(-\mu\Delta x)$ is the probability that the Poisson return process generates a single return event during the incremental displacement Δx and therefore, the vehicle travels a distance x back to facility and subsequently travels distance D_0 to the demand. With probability $\exp(-\mu\Delta x)$, the Poisson return process generates no return event, and the vehicle advances toward the demand by distance Δx and subsequently covers the remaining distance $D(x+\Delta x;c)$ to the demand. The term $o(\Delta x)$ represents $(\Delta x)^2$ and higher order terms. Upon simplifying, we get

$$\begin{aligned} &(D(x+\Delta x;c) - D(x;c)) - \mu\Delta x D(x+\Delta x;c) \\ &\qquad = -\mu x\Delta x - (1+\mu D_0)\Delta x + o(\Delta x). \end{aligned}$$

Dividing both sides by Δx and in the limit $\Delta x \to 0^+$, we obtain the following differential equation for D,

$$\begin{aligned} \frac{dD}{dx} - \mu D &= -\mu x - (1+\mu D_0) \\ \implies D &= -\mathbf{e}^{\mu x}\int \mathbf{e}^{-\mu x}(\mu x + (1+\mu D_0))dx \\ &= \mathbf{e}^{\mu x}(x\mathbf{e}^{-\mu x} + \frac{1}{\mu}\mathbf{e}^{-\mu x} + \frac{1+\mu D_0}{\mu}\mathbf{e}^{-\mu x} + K) \\ &= x + \frac{1}{\mu} + \frac{1+\mu D_0}{\mu} + K\mathbf{e}^{\mu x}. \end{aligned}$$

Using the boundary condition, $D(c;c) = 0$, we have

$$K = -\mathbf{e}^{-\mu c}(c + \frac{2+\mu D_0}{\mu}).$$

Thus, $D(x;c) = -\mathbf{e}^{-\mu(c-x)}(c + \frac{2}{\mu} + D_0) + x + \frac{2}{\mu} + D_0$.

Substituting $x = 0$, and solving for D_0, we have

$$D_0 = -\mathbf{e}^{-\mu c}(c + \frac{2}{\mu} + D_0) + \frac{2}{\mu} + D_0 \Rightarrow D_0 = \frac{2}{\mu}(\mathbf{e}^{\mu c} - 1) - c.$$

□

Lemma III.1 immediately yields the following result on the expected service time in the light load case.

Theorem III.2 (Light load) *In the limit as $\lambda \to 0^+$, the steady-state service time using the FCFS policy is given by*

$$\mathbb{E}[T] = \bar{s} + \frac{1}{A}\int_{\mathcal{E}} \left(\frac{2}{\mu}(\mathbf{e}^{\mu\sqrt{x^2+y^2}} - 1) - \sqrt{x^2+y^2}\right) dx\, dy.$$

Proof: The proof follows from Lemma III.1 together with the fact that in light load, the number of unserviced demands tends to zero. Therefore, the expected service time for a demand is the expected time taken by the vehicle to reach the demand from the facility. □

Remark III.3 (Optimality) *FCFS is optimal in the light load if the facility is located at the point (*weighted median *[18]) that minimizes the expected distance to a demand, which happens to be the case in the present analysis.*

IV. Heavy Load Analysis

We now present the analysis for the heavy load case, i.e., $\rho \to 1^-$ such that $\lambda \to +\infty$ and $\bar{s} \to 1/\lambda$. We begin with a lower bound on $\mathbb{E}[T]$ which becomes a fundamental limit to the problem. We then present a policy for the vehicle and provide an upper bound on the expected service time.

A. Lower bound on $\mathbb{E}[T]$

We first provide a master lower bound on the travel distance between any two demands.

Lemma IV.1 (Master lower bound on travel distance) *Let i and j denote two demands with d_i as the distance between the demand i and the facility, and d_{ij} being the distance between the two demands. Let $c_{\min}$ be the distance between the facility and the nearest demand. Then, a lower bound on the expected travel distance between demands for a vehicle under the recall process is*

$$\bar{d}_f \;\geq\; \mathbb{E}[\mathbf{e}^{-\mu d_{ij}} d_{ij}] + \mathbb{E}[(1 - \mathbf{e}^{-\mu d_{ij}})(d_i + c_{\min})], \quad (3)$$

where the expectation on the right hand side is with respect to the joint distribution of d_i, d_{ij} and $c_{\min}$.

Proof: Consider the j-th demand and suppose that the vehicle has just completed service of the i-th demand, is at the i-th demand location, and is headed toward the j-th demand. The scenario below yields a lower bound on $\bar{d}_f$:

Conditioned on i and j, with probability $\mathbf{e}^{-\mu d_{ij}}$, the vehicle reaches demand j along the shortest path from demand i to j and therefore, covers a distance of d_{ij}. If the vehicle is recalled, which happens with probability $1 - \mathbf{e}^{-\mu d_{ij}}$, then it goes to the facility and then from the facility to the demand *nearest to the facility*, as illustrated in Figure 1.

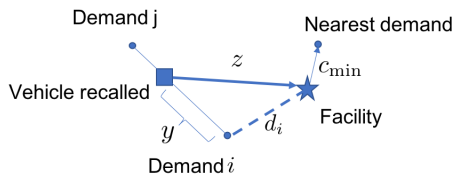

Fig. 1. A snapshot of a recall event enroute demand j from demand i. The vehicle moves to the facility and then to the nearest demand from the facility in the best case scenario.

Applying the triangle inequality (cf. Figure 1), we conclude that the total distance travelled, $y+z+c_{\min} \geq d_i + c_{\min}$. Thus, conditioned on the locations of demands i and j, the expected distance (with respect to the recall process) covered by the vehicle is at least

$$\mathbf{e}^{-\mu d_{ij}} d_{ij} + (1 - \mathbf{e}^{-\mu d_{ij}})(d_i + c_{\min}).$$

Unconditioning on $c_{\min}$ and the locations of demands i and j, we obtain the claim. □

We now present a heavy load lower bound.

Theorem IV.2 (Heavy load lower bound) *Suppose that the demands are generated uniformly randomly in $\mathcal{E}$ and Poisson in time with rate $\lambda \to +\infty$ and that the recall is modeled as a Poisson process with rate $\mu \geq 0$. Then, for any stable policy, in the limit of $\rho \to 1^-$,*

$$\mathbb{E}[T] \geq \frac{\beta_{\text{TSP}}^2}{2} \mathbf{e}^{-2\mu\sqrt{2A}} \frac{\lambda A}{(1-\rho)^2} + \bar{s}.$$

Proof: This proof involves steps that follow the proof of Theorem 2 from [19], but leverage an improved lower bound on the expected distance between demands from Theorem II.2. Since $\mathbb{E}[T]$ is bounded by assumption, let the expected number of demands in the environment be denoted by $\mathbb{E}[n]$. From Lemma IV.1, under the recall process,

$$\bar{d}_f \geq \mathbb{E}[d\mathbf{e}^{-\mu d}],$$

where d is the distance between two demands in the heavy load regime and the expectation is taken over the distribution of the demands. Note that we keep only the first term on the right hand side from Lemma IV.1, and neglect the other two non-negative terms. If the density of the random variable d is denoted by $f_d(\cdot)$, then note that $0 \leq d \leq \sqrt{2A}$. Therefore,

$$\begin{aligned}\bar{d}_f \geq \mathbb{E}[d\mathbf{e}^{-\mu d}] &= \int_0^{\sqrt{2A}} x\mathbf{e}^{-\mu x} f_d(x)\, dx \\ &\geq \mathbf{e}^{-\mu\sqrt{2A}} \int_0^{\sqrt{2A}} x f_d(x)\, dx \\ &= \mathbf{e}^{-\mu\sqrt{2A}} \mathbb{E}[d] \geq \mathbf{e}^{-\mu\sqrt{2A}} \frac{\beta_{\text{TSP}}}{\sqrt{2}} \sqrt{\frac{A}{\mathbb{E}[n]}},\end{aligned}$$

where we used the fact that $\mathbf{e}^{-\mu x}$ is monotonically decreasing with x and the final step follows by applying Theorem II.2. The system can now be viewed as a *queue* in which the service time of a demand is replaced by a sum of the on-site service time and the time taken to reach the demand [19]. Using a necessary condition for stability of a queue, i.e., the expected time taken to service a demand should not exceed the average time between demand arrivals (cf. [19, Proof of Theorem 2]),

$$\bar{s} + \bar{d}_f \leq \frac{1}{\lambda} \;\Leftrightarrow\; \mathbf{e}^{-\mu\sqrt{2A}} \frac{\beta_{\text{TSP}}}{\sqrt{2}} \sqrt{\frac{A}{\mathbb{E}[n]}} \leq \frac{1}{\lambda} - \bar{s},$$

and from the relation, $\mathbb{E}[n] = \lambda(\mathbb{E}[T] - \bar{s})$, we obtain

$$\mathbf{e}^{-\mu\sqrt{2A}} \frac{\beta_{\text{TSP}}}{\sqrt{2}} \sqrt{\frac{A}{\lambda(\mathbb{E}[T] - \bar{s})}} \leq \frac{1}{\lambda} - \bar{s},$$

which, upon rearranging, yields the desired claim. □

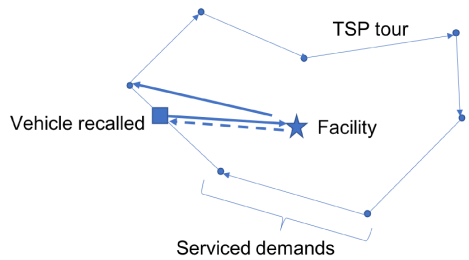

Fig. 2. A snapshot of an epoch using the TSP-based policy. As per the policy, if the vehicle is recalled, then it moves to the facility and then directly on to the next demand in the tour (represented by the arrows). For the purpose of analysis, we analyze the variant in which after moving to the facility, the vehicle moves to the point of recall on the TSP tour and then continues following the tour (represented by the dashed arrow).

B. TSP-based Policy for Heavy Load

This section presents and analyzes a TSP-based policy for the heavy load under the vehicle recall model considered. The policy is an extension of the TSP-based policy from [19] and is described as follows:

Consider the k-th epoch in which there are n_k unserviced demands in $\mathcal{E}$. The vehicle computes a TSP tour through these n_k demands and services them as per the order in the tour. If the vehicle gets recalled, it returns to the facility and then continues the service of the subsequent demand in the tour. Upon completion of service of *all* of the n_k demands in the k-th epoch, the vehicle repeats this procedure at the $k+1$-th epoch.

The following guarantee holds for the TSP-based policy.

Theorem IV.3 (TSP-based policy) *Suppose that the demands are generated uniformly randomly in $\mathcal{E}$ and Poisson in time with rate $\lambda \to +\infty$ and that the recall is modeled as a Poisson process with rate $\mu \geq 0$. Then, in the limit of $\rho \to 1^-$,*

$$\mathbb{E}[T_{\text{TSP}}] \leq \beta_{\text{TSP}}^2 (2\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1)^2 \frac{A\lambda}{(1-\rho)^2} + \bar{s}.$$

Proof: Let Q_k denote the set of demands in the k-th epoch with $|Q_k| = n_k$. For the purpose of analysis, we consider the following modification of the TSP-based policy, as illustrated in Figure 2: whenever the vehicle is recalled, it returns to location from which it was recalled before resuming the tour. Suppose that the recall point is at a distance of r from the

facility. From Lemma III.1, the expected distance (with respect to the recall process) covered by the vehicle to go from the location from which the vehicle was recalled to the facility and back to the recall location is upper bounded by

$$r + \frac{2}{\mu}(\mathbf{e}^{\mu r} - 1) - r \leq \frac{2}{\mu}(\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1),$$

in which we used the fact that $r \leq \sqrt{A/2}$. Due to unit speed of the vehicle, the time taken by the vehicle to complete epoch k using this modification is an upper bound on the time taken by the TSP-based policy.

If the number of recalls in the k-th epoch are m_k, then the time taken T_k to complete the tour in the k-th epoch satisfies

$$T_k \leq \mathrm{ETSP}(Q_k) + m_k \frac{2}{\mu}(\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1) + n_k \bar{s}.$$

Since the recall process is Poisson distributed with rate μ and is independent of the demand generation process, conditioning on n_k and taking expectation with respect to the recall process,

$$\begin{aligned}
\mathbb{E}[T_k \mid n_k] &\leq \mathrm{ETSP}(Q_k) + \frac{2}{\mu}(\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1)\mathbb{E}[m_k] + n_k\bar{s} \\
&\leq \mathrm{ETSP}(Q_k) + \frac{2}{\mu}(\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1)\mu\,\mathrm{ETSP}(Q_k) + n_k\bar{s} \\
&= \mathrm{ETSP}(Q_k)(2\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1) + n_k\bar{s} \\
&= \beta_{\mathrm{TSP}}\sqrt{An_k}(2\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1) + n_k\bar{s},
\end{aligned}$$

where we applied Theorem II.1 at the final step. Unconditioning on n_k, we obtain

$$\begin{aligned}
\mathbb{E}[T_k] &\leq \beta_{\mathrm{TSP}}(2\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1)\sqrt{A}\mathbb{E}[\sqrt{n_k}] + \mathbb{E}[n_k]\bar{s} \\
&\leq \beta_{\mathrm{TSP}}(2\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1)\sqrt{A}\sqrt{\mathbb{E}[n_k]} + \mathbb{E}[n_k]\bar{s},
\end{aligned}$$

where the second inequality is obtained by applying Jensen's inequality to $\sqrt{n_k}$ [7]. Steady-state is achieved by this policy if $\mathbb{E}[n_k] = \lambda(\mathbb{E}[T_k] - \bar{s})$. Substituting this condition into the above equation,

$$\begin{aligned}
\mathbb{E}[T_k] - \lambda\bar{s}(\mathbb{E}[T_k] - \bar{s}) &\leq \beta_{\mathrm{TSP}}(2\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1)\sqrt{A\lambda(\mathbb{E}[T_k] - \bar{s})} \\
\Rightarrow (\mathbb{E}[T_k] - \bar{s})(1 - \lambda\bar{s}) &\leq \beta_{\mathrm{TSP}}(2\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1)\sqrt{A\lambda(\mathbb{E}[T_k] - \bar{s})} \\
\Rightarrow \sqrt{\mathbb{E}[T_k] - \bar{s}} &\leq \beta_{\mathrm{TSP}}(2\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1)\frac{\sqrt{A}}{(1-\rho)},
\end{aligned}$$

where in the second step, we subtracted $\bar{s}$, a positive number only from the left hand side. This yields the claim. □

Remark IV.4 (Factor of optimality) *From Theorems IV.2 and IV.3, for a fixed μ, the TSP-based policy achieves a factor of optimality in terms of A and ρ given by*

$$2\frac{(2\mathbf{e}^{\mu\sqrt{\frac{A}{2}}} - 1)^2}{\mathbf{e}^{-2\mu\sqrt{2A}}},$$

which tends to 2 in the limit of $\mu \to 0^+$. This expression suggests that if μ were to scale faster than $O(1/\sqrt{A})$, then it would mean that the vehicle would be recalled frequently enough that it spends most of its time simply moving back and forth between the facility and a demand location, leading to poor performance in terms of expected service time. On the other hand, if $\mu = K/\sqrt{A}$, we obtain a constant factor *approximation for this problem with the factor given by*

$$2\frac{(2\mathbf{e}^{K\sqrt{\frac{1}{2}}} - 1)^2}{\mathbf{e}^{-2\sqrt{2}K}}.$$

This constant is below 4 for a value of $K \leq 1/10$.

V. Numerical Studies

We now present results of numerical studies in order to quantify performance of the TSP-based policy in practice to assess the gap with the theory and to study the performance in the non-asymptotic regimes of the problem parameters. The `linkern`[1] solver was used to generate approximations to the ETSP at every iteration of the policy.

The numerical implementations were performed as follows: with an initial number of demands given by the arrival rate λ times the lower bound on $\mathbb{E}[T]$, we run the TSP-based policy for a sufficiently large number of epochs (in our case, 20 epochs) to ensure steady-state is reached. Then, we report the average of the service times at steady-state. We assume an environment with area $A = 10$ and mean onsite service time $\bar{s} = 0.1$. Figure 3 summarizes the expected service time at steady-state for varying values of the demand arrival rate λ with the value of $\mu = 0.1$. The results suggest that although Theorems IV.2 and IV.3 have been proven to hold only in limiting regimes of $\lambda \to +\infty$ and $\rho \to 1^-$, the bounds hold empirically for non-asymptotic regimes as well.

We also numerically study a simple *nearest neighbor (NN)* policy in which the vehicle services the demand nearest to the most recently serviced demand. This policy is difficult to analyze, but it is computationally more efficient than the TSP-based policy. However, the TSP-based policy has a much lower expected service time than the NN. Further, there exist demands that do not get served over the duration of the simulation in the NN policy, thereby leading to $\mathbb{E}[T_{\mathrm{NN}}] = +\infty$. In Figure 3, we have reported only the average service times for the demands that have been served in the duration of the simulation.

Next, we empirically study the performance under a different recall model, instead of the Poisson process analyzed in this work. This is motivated by modeling of failure events that make the vehicle return to the facility. We consider a more general, linear hazard function [22], $h(x) = \mu + bx$, which leads to $F(x) = 1 - \mathbf{e}^{-\mu x - bx^2/2}$, as the cumulative distribution function for the recall. Such increasing hazard rate captures the fact that the recall probability increases with time. Here, μ is a base recall rate and b is the rate of increase of the recall rate. Figure 4 summarizes the results of the TSP-based policy applied to this model and suggests that the expected service time: 1) varies exponentially with λ and 2) is more sensitive to b than to μ.

[1]The TSP solver `linkern` is freely available for academic research use at `http://www.math.uwaterloo.ca/tsp/concorde/`.

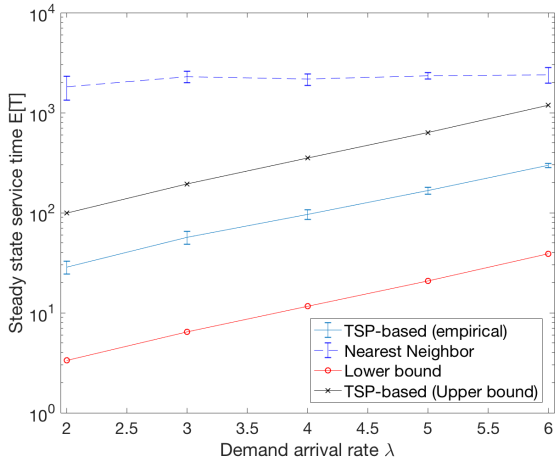

Fig. 3. Performance of the TSP-based and the NN policies. For the NN policy, this plot reports only the average service times for the demands that have been served, since there exist demands that do not get served over the duration of the simulation.

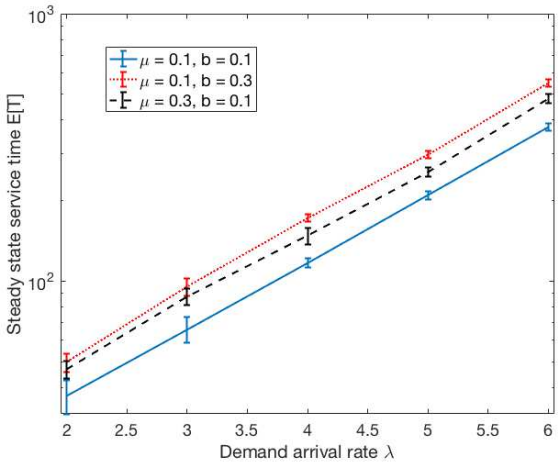

Fig. 4. Empirical performance of the TSP-based policy for different values of the parameters μ, b of a linear hazard function whose cumulative distribution function is given by $F(t) = 1 - \mathbf{e}^{-\mu t - bt^2/2}$.

VI. Conclusion and Future Directions

This paper introduced a new variant of DVR problems in which a single vehicle gets recalled to a central facility in the environment as per a Poisson process. The goal was to design service policies to minimize the steady-state expected service time of demands that appear in the environment via another Poisson process. We analyzed an FCFS policy in light load. In heavy load, we derived a policy independent lower bound on the expected service time as a function of the problem parameters. We then designed a service policy by adapting a well-known policy based on repeated computation of the ETSP tour through unserviced demands. We provided an upper bound on the steady-state service time for this policy that performs within a constant factor of the lower bound in terms of the demand arrival rate. Finally, we presented a numerical study in non-asymptotic parameter regimes, comparison with nearest neighbor policy and studied a recall model beyond the scope of the presented analysis.

In future, we plan to explore theoretical developments extending this work to more general failure models with memory, such as the one with linear hazard rate. Also of interest is the version of this problem with multiple vehicles.

References

[1] G. A. Boyce, "The rugged border: Surveillance, policing and the dynamic materiality of the us/mexico frontier," *Environment and Planning D: Society and Space*, vol. 34, no. 2, pp. 245–262, 2016.

[2] S. Poikonen, X. Wang, and B. Golden, "The vehicle routing problem with drones: Extended models and connections," *Networks*, vol. 70, no. 1, pp. 34–43, 2017.

[3] P. Toth and D. Vigo, "An overview of vehicle routing problems," in *The vehicle routing problem*. SIAM, 2002, pp. 1–26.

[4] H. N. Psaraftis, "Dynamic vehicle routing problems," *Vehicle routing: Methods and studies*, vol. 16, pp. 223–248, 1988.

[5] D. J. Bertsimas and G. Van Ryzin, "Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles," *Operations Research*, vol. 41, no. 1, pp. 60–76, 1993.

[6] J. D. Papastavrou, "A stochastic and dynamic routing policy using branching processes with state dependent immigration," *European Journal of Operational Research*, vol. 95, no. 1, pp. 167–177, 1996.

[7] M. Pavone, E. Frazzoli, and F. Bullo, "Adaptive and distributed algorithms for vehicle routing in a stochastic and dynamic environment," *IEEE Trans. on Auto. Cont.*, vol. 56, no. 6, pp. 1259–1274, 2011.

[8] S. L. Smith, M. Pavone, F. Bullo, and E. Frazzoli, "Dynamic vehicle routing with priority classes of stochastic demands," *SIAM Journal on Control and Optimization*, vol. 48, no. 5, pp. 3224–3245, 2010.

[9] H. A. Waisanen, D. Shah, and M. A. Dahleh, "A dynamic pickup and delivery problem in mobile networks under information constraints," *IEEE Trans. on Auto. Cont.*, vol. 53, no. 6, pp. 1419–1433, 2008.

[10] K. Treleaven, M. Pavone, and E. Frazzoli, "Asymptotically optimal algorithms for one-to-one pickup and delivery problems with applications to transportation systems," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2261–2276, 2013.

[11] K. Savla, E. Frazzoli, and F. Bullo, "Traveling salesperson problems for the Dubins vehicle," *IEEE Transactions on Automatic Control*, vol. 53, no. 6, pp. 1378–1391, 2008.

[12] K. Savla, F. Bullo, and E. Frazzoli, "Traveling salesperson problems for a double integrator," *IEEE Transactions on Automatic Control*, vol. 54, no. 4, pp. 788–793, 2009.

[13] S. Itani, E. Frazzoli, and M. A. Dahleh, "Dynamic traveling repairperson problem for dynamic systems," in *2008 47th IEEE Conference on Decision and Control*. IEEE, 2008, pp. 465–470.

[14] A. Arsie, K. Savla, and E. Frazzoli, "Efficient routing algorithms for multiple vehicles with no explicit communications," *IEEE Transactions on Automatic Control*, vol. 54, no. 10, pp. 2302–2317, 2009.

[15] 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.

[16] S. D. Bopardikar, S. L. Smith, and F. Bullo, "On dynamic vehicle routing with time constraints," *IEEE Transactions on Robotics*, vol. 30, no. 6, pp. 1524–1532, 2014.

[17] P. Agharkar, S. D. Bopardikar, and F. Bullo, "Vehicle routing algorithms for radially escaping targets," *SIAM Journal on Control and Optimization*, vol. 53, no. 5, pp. 2934–2954, 2015.

[18] 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.

[19] D. J. Bertsimas and G. Van Ryzin, "A stochastic and dynamic vehicle routing problem in the Euclidean plane," *Operations Research*, vol. 39, no. 4, pp. 601–615, 1991.

[20] J. Beardwood, J. H. Halton, and J. M. Hammersley, "The shortest path through many points," in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 55, no. 4. Cambridge University Press, 1959, pp. 299–327.

[21] H. Xu, "Optimal policies for stochastic and dynamic vehicle routing problems," Ph.D. dissertation, Mass. Inst. of Technology, 1994.

[22] H. Rinne, "The Hazard Rate: Theory and Inference," Justus Liebig University, Tech. Rep. [Online]. Available: http://geb.uni-giessen.de/geb/volltexte/2014/10793/pdf/RinneHorst_hazardrate_2014.pdf