ACC 2009, Submitted
St. Louis, MO

Dynamic Vehicle Routing with Moving Demands – Part I: Low speed demands and high arrival rates

Shaunak D. Bopardikar Stephen L. Smith Francesco Bullo João P. Hespanha

***Abstract*— We introduce a dynamic vehicle routing problem in which demands arrive via a temporal Poisson process with a certain arrival rate, and uniformly distributed along a line segment. Upon arrival, the demands move in a fixed direction perpendicular to the line with a fixed speed. A service vehicle, modeled as a first-order integrator with speed greater than that of the demands, seeks to serve these mobile demands. For the existence of any stabilizing service policy, we determine a necessary condition on the arrival rate of the demands in terms of the problem parameters; (i) the speed ratio between the demand and service vehicle, and (ii) the length of the line segment on which demands arrive. Next, we propose a novel service policy for the vehicle that involves servicing the outstanding demands as per the traveling salesperson path (t-TSP) through the moving demands. We derive a sufficient condition on the arrival rate of the demands for stability of the TSP-based policy, in terms of the problem parameters. We show that in the limiting case in which the demands move much slower than the service vehicle, the necessary and the sufficient conditions on the arrival rate are within a constant factor. We also provide an upper bound on the steady-state expected time spent by each demand before being served.**

I. INTRODUCTION

Dynamic vehicle routing problems such as the dynamic traveling repairperson problem (DTRP), consider one (or more) service vehicles that seek to serve demands that arrive via some spatio-temporal process in a region, and upon arrival the demands remain at their location until they are served. In this work, we introduce a dynamic vehicle routing problem in which the demands move with a specified velocity upon arrival, and we design policies for a single vehicle that seeks to serve them. This problem has applications in areas such as perimeter defense, wherein the demands could be visualized as moving targets trying to cross a region under surveillance by a UAV. Another application is in the automation industry where the demands are objects that arrive continuously on a conveyor belt and a robotic arm seeks to perform a pick-and-place operation on them.

The DTRP was first introduced in [1] in which the goal is to minimize the expected time spent by each demand before being served. In [1] the authors propose a policy that is optimal in the case of low arrival rate, and several policies within a constant factor of the optimal in the case of high arrival rate. In [2], they also study multiple service vehicles, and vehicles with finite service capacity. In [3], a single policy is proposed 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. In [4], decentralized policies are developed for the multiple service vehicle versions of the DTRP.

The authors are with the Center for Control, Dynamical Systems and Computation, University of California at Santa Barbara, Santa Barbara, CA 93106, USA, {shaunak,stephen,bullo}@engineering.ucsb.edu, hespanha@ece.ucsb.edu

The Euclidean traveling salesperson problem (ETSP) consists of determining the minimum length tour through a given set of points in a region [5]. Vehicle routing with targets moving on straight lines was introduced in [6], where a fixed number of targets move in the same direction with fixed speed, and the problem is to catch the maximum number of targets before they cross a finish line. A variation of this problem with target motion on piece-wise straight line paths and with varying target speeds has been addressed in [7]. For the case in which there is no finish line, termed as the translational traveling salesperson problem (t-TSP), a polynomial-time approximation scheme has been proposed in [8] to catch all targets in minimum time. Other variants of the ETSP in which the points are allowed to move in different directions have been addressed in [8] and in [9].

We introduce a dynamic vehicle routing problem in which demands arrive via a temporal Poisson process with rate λ, and uniformly randomly on a line segment of finite length W. Upon arrival, the demands move in a fixed direction perpendicular to the line and with a fixed speed $v < 1$. A service vehicle, modeled as a first-order integrator with unit speed, seeks to serve these mobile demands.

Our main contributions are as follows: First, we show that to ensure the existence of a stabilizing policy, i.e., a finite expected time spent by a demand in the environment, we must have $\lambda \leq 4/vW$. Second, we propose a novel service policy which involves servicing all of the outstanding demands as per the translational traveling salesperson path (TSP) through them. We show that a sufficient condition for stability of this TSP-based policy is $\lambda < (1-v^2)^{3/2}/2vW(1+v)^2$. For this policy, we also obtain an upper bound on the steady-state expected time a demand spends in the environment before being served. As the arrival rate $\lambda \to +\infty$, the necessary stability condition implies that the demands must have $v \to 0^+$. This regime of low demand speed with high arrival rate is the focus of this paper. In this regime, the necessary and the sufficient stability conditions on the arrival rate mentioned above are within a constant factor.

In companion paper [10], we analyze a first-come-first-served (FCFS) policy in which the demands are served in the order of their arrival. We show that in the regime of $\lambda \to 0^+$, the FCFS policy minimizes the expected time spent by a demand before being served; while in the regime of $v \to 1^-$, the FCFS is the optimal policy. Thus, for low demand speeds the TSP-based policy can stabilize higher arrival rates, while for high demand speeds the FCFS can stabilize higher arrival rates. This is summarized in Fig. 1.

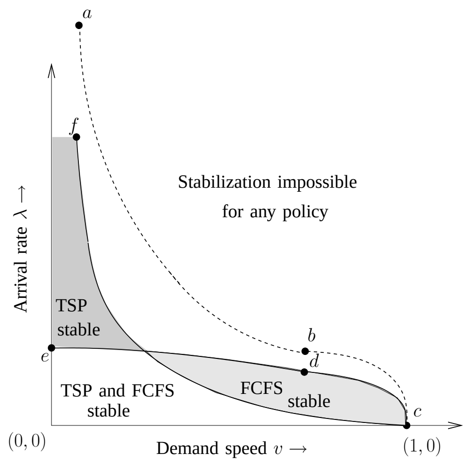

Fig. 1. A summary of stability regions for the TSP-based policy and the FCFS policy. Stable service policies can exist only for the region under the dotted line. Curve a-b is due to Theorem IV.1, and curve c-f is due to Theorem V.1. Curves b-c, c-d and d-e, are established in [10].

This paper is organized as follows: we begin with background results on the traveling salesperson problems in Section II. The problem formulation is presented in Section III. The necessary condition for stability is derived in Section IV. The TSP-based service policy and the main results are presented in Section V. Simulation results are presented in Section VI.

II. Background results on Traveling Salesperson Problems

In this section we review several results on determining shortest paths through sets of points. In what follows, given a set of points in the plane, we are interested in the length of a path through the points that *is not closed*. These results will be applied in the analysis in Section V.

A. The Euclidean Traveling Salesperson Path (ETSP)

We are interested in the following Euclidean TSP problem.

> Given n static points placed in $\mathbb{R}^2$, determine the length of the shortest path through all the points.

An upper bound on the length of such a path for points in a unit square was given by Few [11]. Here we extend Few's bound to the case of points in a rectangular region. For completeness, we have included the proof in the Appendix.

Lemma II.1 (Euclidean TSP length) *Given n points in a $1 \times h$ rectangle in the plane, where $h \in \mathbb{R}_{>0}$, there exists a path that starts from one of the edges of the rectangle, passes through each of the n points exactly once, and terminates on the opposite edge with length upper bounded by*

$$\sqrt{2hn} + h + 5/2.$$

We will also require a result on the length of a path through a large number of points. Given a set $\mathcal{Q}$ of n points in $\mathbb{R}^2$, let $\mathrm{ETSP}(\mathcal{Q})$ denote the minimum length of a path through all the points in $\mathcal{Q}$. The following is an established result.

Theorem II.2 (Asymptotic TSP length, [12]) *If n points are distributed independently and identically uniform in a compact region of area A, then there exists a constant β_{TSP} such that*

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

The current best estimate of the constant is $\beta_{\mathrm{TSP}} \simeq 0.7120$, [13].

B. The Translational Traveling Salesperson Path

Next, we describe the translational TSP which was proposed and solved in [8]. This problem is posed as follows.

> Given a start point $\mathbf{s}(t)$, a set of points $\mathcal{Q}(t) := \{\mathbf{q}_1(t), \dots, \mathbf{q}_n(t)\}$ and a finish point $\mathbf{f}(t)$ all moving with the same constant speed v and in the same direction, determine a salesperson path that starts from $\mathbf{s}$, visits all points in $\mathcal{Q}$ and ends at $\mathbf{f}$, and the length $\mathcal{L}_T(\mathbf{s}, \mathcal{Q}, \mathbf{f})$ of which is minimum.

A solution for this problem is the *Convert-to-Static TSP* method:

(i) Define the map $f_v : \mathbb{R}^2 \to \mathbb{R}^2$ by

$$f_v(x, y) = \Big(\frac{x}{\sqrt{1 - v^2}}, \frac{y}{1 - v^2} \Big).$$

(ii) Compute the static TSP that starts at $f_v(\mathbf{s})$, passes through the set of points given by $\{f_v(\mathbf{q}_1), \dots, f_v(\mathbf{q}_n)\}$ and ends at $f_v(\mathbf{f})$.

For this method, the following result is established.

Lemma II.3 (Translational TSP length, [8]) *The length of the translational TSP is*

$$\mathcal{L}_T(\mathbf{s}, \mathcal{Q}, \mathbf{f}) = \mathcal{L}_E(\mathbf{s}, \mathcal{Q}, \mathbf{f}) + \frac{v(y_f - y_s)}{1 - v^2},$$

where $\mathcal{L}_E(\mathbf{s}, \mathcal{Q}, \mathbf{f})$ denotes the length of the static TSP with starting point $f_v(\mathbf{s}) := f_v(x_s, y_s)$, passing through the set of points $\{f_v(\mathbf{q}_1), \dots, f_v(\mathbf{q}_n)\}$, and ending at $f_v(\mathbf{f}) := f_v(x_f, y_f)$.

In other words, the length of the translational TSP is optimal if and only if the length of the TSP in the corresponding static instance is optimal.

III. Problem Formulation

We consider a single service vehicle that seeks to service mobile demands that arrive via a spatio-temporal process on a line segment with length W along the x-axis, termed the *generator*. The vehicle is modeled as a first-order integrator with speed upper bounded by one. The demands arrive uniformly distributed on the generator via a temporal Poisson process with intensity $\lambda > 0$, and move with constant speed $v < 1$ along the positive y-axis, as shown in Fig. 2. We assume that once the vehicle reaches a demand, the demand is served instantaneously. The vehicle is assumed to have unlimited fuel and demand servicing capacity.

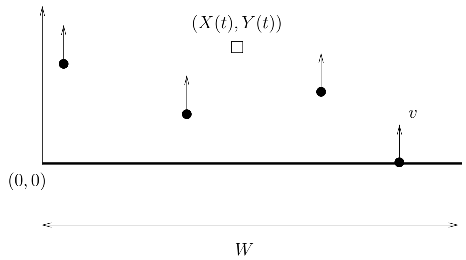

Fig. 2. The problem set-up. The thick line segment is the generator of mobile demands. The dark circle denotes a demand and the square denotes the service vehicle.

We define the environment as $\mathcal{E} := [0, W] \times \mathbb{R}_{\geq 0} \subset \mathbb{R}^2$, and let $\mathbf{p}(t) = [X(t), Y(t)]^T \in \mathcal{E}$ denote the position of the service vehicle at time t. Let $\mathcal{Q}(t) \subset \mathcal{E}$ denote the set of all 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 service vehicle reaches the location of the demand. A static feedback control policy for the system is a map $\mathcal{P} : \mathcal{E} \times 2^{\mathcal{E}} \to \mathbb{R}^2$, assigning a commanded velocity to the service vehicle as a function of the current state of the system: $\dot{\mathbf{p}}(t) = \mathcal{P}(\mathbf{p}(t), \mathcal{Q}(t))$. Let D_i denote the time that the ith demand spends within the set $\mathcal{Q}$, i.e., the delay between the generation of the ith demand and the time it is serviced. The policy $\mathcal{P}$ is *stable* if under its action,

$$\lim_{i \to +\infty} \mathbb{E}\left[D_i\right] < +\infty,$$

i.e., the steady state expected delay is finite. Equivalently, the policy $\mathcal{P}$ is stable if under its action, $\lim_{t \to +\infty} \mathbb{E}\left[n(t)\right] < +\infty$, that is, if the vehicle is able to service demands at a rate that is—on average—at least as fast as the rate at which new demands arrive. In what follows, our goal is to *design stable control policies* for the system.

IV. A Necessary condition for stability

In this section we provide a necessary condition on the arrival rate for the existence of a stabilizing policy. We begin by stating the main result of the section, with the remainder of the section dedicated to its proof.

Theorem IV.1 (Necessary condition for stability) *A necessary condition for the existence of a stabilizing policy is that*

$$\lambda \leq \frac{4}{vW}.$$

Before proving this result we state one of its key consequences.

Corollary IV.2 (Constant fraction service) *A necessary condition for the existence of a policy which services a fraction $c \in (0, 1]$ of the demands is that*

$$\lambda \leq \frac{4}{c^2 vW}.$$

Thus, for a fixed $v < 1$ no policy can service a constant fraction of the demands as $\lambda \to +\infty$.

To prove Theorem IV.1 we begin by looking at the distribution of demands in the service region.

Lemma IV.3 (Poisson point process) *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 $[0, W] \times [0, vt]$ at time t. Then, given a compact region $\mathcal{R}$ of area A contained in $[0, W] \times [0, vt]$,*

$$\mathbb{P}[|\mathcal{R} \cap \mathcal{Q}| = n] = \frac{\mathrm{e}^{-\bar{\lambda} A} (\bar{\lambda} A)^n}{n!}, \quad \textit{where } \bar{\lambda} := \lambda/(vW).$$

Proof: Let $\mathcal{R} = [\ell, \ell + \Delta\ell] \times [h, h + \Delta h]$ be a rectangle contained in $[0, W] \times [0, vt]$ with area $A = \Delta\ell\Delta h$. Let us calculate the probability that at time t, $|\mathcal{R} \cap \mathcal{Q}| = n$ (that is, the probability that $\mathcal{R}$ contains n points in $\mathcal{Q}$). We have

$$\begin{aligned} &\mathbb{P}[|\mathcal{R} \cap \mathcal{Q}| = n] = \\ &\quad \sum_{i=n}^{\infty} \mathbb{P}\left[i \text{ demands arrived in } \left[\frac{h}{v}, \frac{h + \Delta h}{v}\right]\right] \\ &\qquad \times \mathbb{P}[n \text{ of } i \text{ are generated in } [\ell, \ell + \Delta\ell]]. \end{aligned}$$

Since the generation process is temporally Poisson and spatially uniform the above equation can be rewritten as

$$\begin{aligned} \mathbb{P}[|\mathcal{R} \cap \mathcal{Q}| = n] &= \sum_{i=n}^{\infty} \mathbb{P}\left[i \text{ demands arrived in } [0, \Delta h/v]\right] \\ &\quad \times \mathbb{P}[n \text{ of } i \text{ are generated in } [0, \Delta\ell]]. \end{aligned} \tag{2}$$

Now,

$$\mathbb{P}\left[i \text{ demands arrived in } [0, \Delta h/v]\right] = \frac{\mathrm{e}^{-\lambda\Delta h/v} (\lambda\Delta h/v)^i}{i!},$$

and,

$$\mathbb{P}[n \text{ of } i \text{ are in } [0, \Delta\ell]] = \binom{i}{n} \left(\frac{\Delta\ell}{W}\right)^n \left(1 - \frac{\Delta\ell}{W}\right)^{i-n}.$$

So, letting $L := \Delta\ell/W$ and $H := \Delta h/v$, and substituting in the above expressions, Eq. (2) becomes

$$\mathbb{P}[|\mathcal{R} \cap \mathcal{Q}| = n] = \mathrm{e}^{-\lambda H} L^n \sum_{i=n}^{\infty} \frac{(\lambda H)^i}{i!} \binom{i}{n} (1 - L)^{i-n}. \tag{3}$$

Rewriting $(\lambda H)^i$ as $(\lambda H)^n (\lambda H)^{n-i}$, and using the definition

$$\binom{i}{n} = \frac{i!}{n!(i-n)!},$$

we can write Eq. (3) as

$$\begin{aligned} \mathbb{P}[|\mathcal{R} \cap \mathcal{Q}| = n] &= \mathrm{e}^{-\lambda H} \frac{(\lambda L H)^n}{n!} \sum_{j=0}^{\infty} \frac{(\lambda H(1 - L))^j}{j!} \\ &= \mathrm{e}^{-\lambda H + \lambda H(1-L)} \frac{(\lambda L H)^n}{n!} \\ &= \mathrm{e}^{-\lambda L H} \frac{(\lambda L H)^n}{n!}. \end{aligned}$$

Finally, since $LH = A/(vW)$, we obtain

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

where $\bar{\lambda} := \lambda/(vW)$. Thus, the result is established for rectangles. However, every compact region can be written as a countable union of rectangles, and thus the result holds for every compact, measurable region contained in $[0, W] \times [0, vt]$. ■

Remark IV.4 (Uniformly distributed demands) Lemma IV.3 shows us that the number of demands in an unserviced region is Poisson distributed with rate $\lambda/(vW)$, and conditioned on this number, the demands are distributed uniformly. □

We now establish a result on the expected time to travel from a demand to its nearest neighbor. For this we require a result on catching a demand in minimum time (cf. Fig. 3).

Proposition IV.5 (Minimum time control, [14]) *Given the locations $\mathbf{p} := (X, Y) \in \mathcal{E}$ and $\mathbf{q} := (x, y) \in \mathcal{E}$ at time t of the vehicle and a demand, respectively, then the motion of the vehicle towards the point $(x, y + vT)$, where*

$$T(\mathbf{p}, \mathbf{q}) := \frac{\sqrt{(1 - v^2)(X - x)^2 + (Y - y)^2}}{1 - v^2} - \frac{v(Y - y)}{1 - v^2},$$

minimizes the time taken by the vehicle to reach the demand.

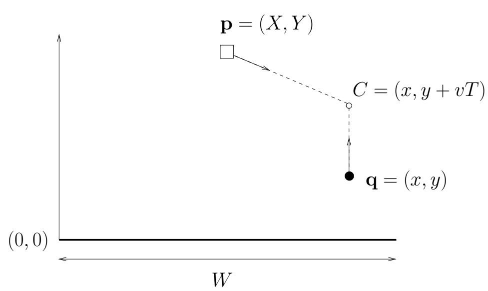

Fig. 3. Illustration of Proposition IV.5.

Lemma IV.6 (Travel time bound) *Consider the set $\mathcal{Q}$ of demands in $\mathcal{E}$ at time t. Let T_d be a random variable giving the minimum amount of time required to travel to a demand in $\mathcal{Q}$ from a vehicle position (X, Y), selected* a priori. *Then*

$$\mathbb{E}[T_d] \geq \frac{1}{2}\sqrt{\frac{vW}{\lambda}}.$$

Proof: To obtain a lower bound on the minimum travel time we can assume that $\mathcal{Q}$ contains many demands (i.e., t is very large), and no demands have been serviced. Consider a demand in $\mathcal{Q}$ with position (x, y) at time t. Using Proposition IV.5, we can write the travel time T from $\mathbf{p} := (X, Y)$ to $\mathbf{q} := (x, y)$ implicitly as

$$T(\mathbf{p}, \mathbf{q})^2 = (X - x)^2 + ((Y - y) - vT(\mathbf{p}, \mathbf{q}))^2. \quad (4)$$

Thus, we can define the set S_T, such that any demand in S_T can be reached from (X, Y) in T time units. From Eq. (4) we see that the set S_T is a circle of radius T centered at $X, Y - vT$. That is,

$$S_T := \{(x, y) \in \mathcal{E} : (X - x)^2 + ((Y - vT) - y)^2 \leq T^2\},$$

where we have omitted T's dependence on $\mathbf{p}$ and $\mathbf{q}$. If the set S_T does not intersect a boundary of $\mathcal{E}$ it has area πT^2, but in general its area is $|S_T| \leq \pi T^2$. Now, by Lemma IV.3 the demands in an unserviced region are uniformly randomly distributed with density $\bar{\lambda} = \lambda/(vW)$. Let us compute the distribution of $T_d := \min_{\mathbf{q} \in \mathcal{Q}} T(\mathbf{p}, \mathbf{q})$. For every vehicle position $\mathbf{p}$ chosen before the generation of demands, the probability that $T_d > T$ is given by

$$\mathbb{P}[T_d > T] = \mathbb{P}[|S_T \cap \mathcal{Q}| = 0] \geq \mathrm{e}^{-\bar{\lambda}|S_T|} \geq \mathrm{e}^{-\lambda\pi T^2/(vW)}.$$

Hence we have

$$\begin{aligned}\mathbb{E}[T_d] &\geq \int_0^{+\infty} \mathbb{P}[T_d > T]dT \geq \int_0^{+\infty} \mathrm{e}^{-\lambda\pi T^2/(vW)} dT \\ &= \frac{\sqrt{\pi}}{2\sqrt{\lambda\pi/(vW)}} = \frac{1}{2}\sqrt{\frac{vW}{\lambda}}.\end{aligned}$$

■

We can now prove Theorem IV.1.

Proof: [Proof of Theorem IV.1] A necessary condition for the stability of any policy (see, for example [1]) is that

$$\lambda\mathbb{E}[T] \leq 1,$$

where $\mathbb{E}[T]$ is the steady-state expected travel time between demands i and $i + 1$. For every policy $\mathbb{E}[T] \geq \mathbb{E}[T_d] \geq \frac{1}{2}\sqrt{\frac{vW}{\lambda}}$. Thus a necessary condition for stability is that

$$\lambda\frac{1}{2}\sqrt{\frac{vW}{\lambda}} \leq 1 \quad \Leftrightarrow \quad \lambda \leq \frac{4}{vW}.$$

■

Finally, Corollary IV.2 follows since in order to service a fraction c we require that $c\lambda\mathbb{E}[T_d] < 1$.

V. The TSP-based policy and the Main result

In this section, we present a novel service policy for the vehicle which is based on computation of the translational traveling salesperson path (TSP) path through successive groups of outstanding demands.

A. The TSP-Based Service Policy and the Main Result

The TSP-based service policy is as follows:

TSP-based service policy

Assumes: Service vehicle has initial position (X, Y), and all demands have lower y-coordinates.

if no outstanding demands in the environment **then**
    Move towards the generating line for a time interval of $Y/(1 + v)$.
else
    Let V be a "virtual" demand located at $(X, 0)$ moving with speed v in the positive y-direction.
    Service all the outstanding demands by following a translational TSP starting from (X, Y), and terminating at virtual demand V.
Repeat.

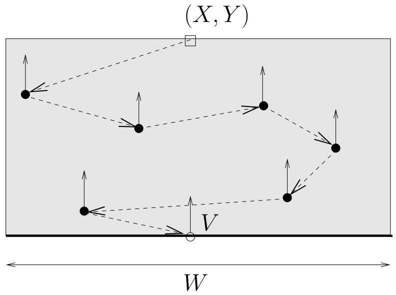

Fig. 4. The TSP-based policy. The vehicle serves all outstanding demands inside the shaded rectangular region $\mathcal{R}(X, Y)$ as per the translational TSP that begins at (X, Y) and terminates at the virtual demand V.

An instance of this policy is illustrated in Fig. 4. The TSP-based service policy gives the following result.

Theorem V.1 (TSP-based policy) (i) *The TSP-based policy is stable if*

$$\lambda < \frac{(1-v^2)^{3/2}}{2vW(1+v)^2}, \text{ and,}$$

(ii) *assuming that the TSP-based policy is stable, the steady state expected time spent by a demand in the environment is no larger than*

$$\frac{5W}{2\sqrt{1-v^2}}\left(\frac{1}{1/(1+v) - \sqrt{2Wv\lambda/(1-v^2)^{3/2}}}\right).$$

Proof: Let $\mathcal{R}(X, Y)$ denote the region $[0, W] \times [0, Y]$ defined by the position (X, Y) of the service vehicle, as shown in Fig. 4. Observe that at the end of every iteration of this policy, all outstanding demands have their y-coordinates less than or equal to that of the vehicle, and hence would be contained in $\mathcal{R}(X, Y)$. Let the vehicle be located at $\mathbf{p}(t_i) = (X(t_i), Y(t_i))$ at time instant t_i. If there are no outstanding demands in $\mathcal{R}(X(t_i), Y(t_i))$, then $\frac{Y(t_i)}{1+v}$ is the distance that the vehicle moves towards the generator. Thus, we have

$$Y(t_{i+1}) = Y(t_i) - \frac{Y(t_i)}{1+v} = \frac{vY(t_i)}{1+v},$$

if there are no unserviced demands in $\mathcal{R}(X(t_i), Y(t_i))$ at time t_i. Otherwise, if there exist unserviced demands $\{\mathbf{q}_1, \ldots, \mathbf{q}_{n_i}\}$ where $n_i \geq 1$, in $\mathcal{R}(X(t_i), Y(t_i))$, then we have

$$Y(t_{i+1}) = v\mathcal{L}_T(\mathbf{p}(t_i), \{\mathbf{q}_1, \ldots, \mathbf{q}_{n_i}\}, V(t_i)),$$

where $\mathcal{L}_T(\mathbf{p}(t_i), \{\mathbf{q}_1, \ldots, \mathbf{q}_{n_i}\}, V(t_i))$ is the time taken for the vehicle as per the translational TSP that begins at $\mathbf{p}(t_i)$, serves all n_i demands and ends at the virtual demand $V(t_i)$. Since the distribution of the demands inside $\mathcal{R}(X(t_i), Y(t_i))$ is spatially Poisson (cf. Lemma IV.3 from Section IV), we have

$$\begin{aligned}
Y(t_{i+1}) &= \frac{vY(t_i)}{1+v}, \quad \text{w.p. } \mathrm{e}^{-\bar{\lambda}A},\\
&= v\mathcal{L}_T(\mathbf{p}(t_i), \{\mathbf{q}_1\}, V(t_i)), \quad \text{w.p. } (\bar{\lambda}A)\mathrm{e}^{-\bar{\lambda}A},\\
&= v\mathcal{L}_T(\mathbf{p}(t_i), \{\mathbf{q}_1, \mathbf{q}_2\}, V(t_i)), \quad \text{w.p. } \frac{(\bar{\lambda}A)^2}{2!}\mathrm{e}^{-\bar{\lambda}A},
\end{aligned}$$

and so on, where $A = WY(t_i)$ is the area of $\mathcal{R}(X(t_i), Y(t_i))$. We now seek an upper bound for the length $\mathcal{L}_T(\mathbf{p}(t_i), \{\mathbf{q}_1, \ldots, \mathbf{q}_{n_i}\}, V(t_i))$ of the translational TSP for which we use the Convert-to-Static TSP method (cf. Section II-A). For $n_i = k \geq 1$, invoking Lemma II.3 and writing $Y_i := Y(t_i)$, for convenience,

$$\begin{aligned}
&\mathcal{L}_T(\mathbf{p}(t_i), \{\mathbf{q}_1, \ldots, \mathbf{q}_k\}, V(t_i))\\
&= \mathcal{L}_E(\mathbf{p}(t_i), \{\mathbf{q}_1, \ldots, \mathbf{q}_k\}, V(t_i)) + \frac{v(y_{V(t_i)} - Y_i)}{1-v^2}\\
&= \mathcal{L}_E(\mathbf{p}(t_i), \{\mathbf{q}_1, \ldots, \mathbf{q}_k\}, V(t_i)) - \frac{vY_i}{1-v^2}\\
&\leq \sqrt{\frac{2WY_ik}{(1-v^2)^{3/2}}} + \frac{Y_i}{1+v} + \frac{5W}{2\sqrt{1-v^2}},
\end{aligned}$$

where the second equality is due to $y_{V(t_i)} = 0$, and the inequality is obtained using Lemma II.1. Thus, we have

$$\begin{aligned}
&\mathbb{E}\left[Y_{i+1}|Y_i\right] \leq v\frac{Y_i}{1+v}\mathrm{e}^{-\bar{\lambda}A} +\\
&v\sum_{k=1}^{\infty}\left(\sqrt{\frac{2WY_ik}{(1-v^2)^{3/2}}} + \frac{Y_i}{1+v} + \frac{5W}{2\sqrt{1-v^2}}\right)\frac{(\bar{\lambda}A)^k}{k!}\mathrm{e}^{-\bar{\lambda}A},
\end{aligned}$$

where $\bar{\lambda} = \lambda/vW$ from Lemma IV.3. Collecting the terms with $vY_i/(1+v)$ together, we obtain

$$\begin{aligned}
&\mathbb{E}\left[Y_{i+1}|Y_i\right] \leq \frac{vY_i}{1+v}\sum_{k=0}^{\infty}\frac{(\bar{\lambda}A)^k}{k!}\mathrm{e}^{-\bar{\lambda}A} +\\
&\qquad \sum_{k=1}^{\infty}\left(\sqrt{\frac{2v^2WY_ik}{(1-v^2)^{3/2}}} + \frac{5vW}{2\sqrt{1-v^2}}\right)\frac{(\bar{\lambda}A)^k}{k!}\mathrm{e}^{-\bar{\lambda}A}\\
&= \frac{vY_i}{1+v} + \sqrt{\frac{2v^2W}{(1-v^2)^{3/2}}}\mathbb{E}\left[\sqrt{n_iY_i}|Y_i\right] + \frac{5vW(1-\mathrm{e}^{-\bar{\lambda}A})}{2\sqrt{1-v^2}}\\
&\leq \frac{vY_i}{1+v} + \sqrt{\frac{2v^2W}{(1-v^2)^{3/2}}}\sqrt{Y_i}\mathbb{E}\left[\sqrt{n_i}|Y_i\right] + \frac{5vW}{2\sqrt{1-v^2}}\\
&\leq \frac{vY_i}{1+v} + \sqrt{\frac{2v^2W}{(1-v^2)^{3/2}}}\sqrt{Y_i}\sqrt{\mathbb{E}\left[n_i|Y_i\right]} + \frac{5vW}{2\sqrt{1-v^2}}\\
&= \frac{vY_i}{1+v} + \sqrt{\frac{2v^2W}{(1-v^2)^{3/2}}}\sqrt{Y_i}\sqrt{\frac{\lambda WY_i}{vW}} + \frac{5vW}{2\sqrt{1-v^2}}\\
&= \frac{vY_i}{1+v} + \sqrt{\frac{2v^2W}{(1-v^2)^{3/2}}}\sqrt{\frac{\lambda}{v}}Y_i + \frac{5vW}{2\sqrt{1-v^2}},
\end{aligned}$$

where the inequality in the fourth step follows by applying Jensen's inequality to the conditional expectation and the equality in the fifth step is due to the arrival process being spatially Poisson (cf. Lemma IV.3). Using the law of iterated expectation, we have

$$\begin{aligned}
&\mathbb{E}\left[Y_{i+1}\right] = \mathbb{E}\left[\mathbb{E}\left[Y_{i+1}|Y_i\right]\right]\\
&\leq \frac{v}{1+v}\mathbb{E}\left[Y_i\right] + \sqrt{\frac{2v\lambda W}{(1-v^2)^{3/2}}}\mathbb{E}\left[Y_i\right] + v\frac{5W}{2\sqrt{1-v^2}}, \quad (5)
\end{aligned}$$

which is a linear recurrence in $\mathbb{E}[Y_i]$. Thus, $\lim_{i\to+\infty} \mathbb{E}[Y_i]$ is finite if

$$\frac{v}{1+v} + \sqrt{\frac{2Wv\lambda}{(1-v^2)^{3/2}}} < 1 \Leftrightarrow \lambda < \frac{(1-v^2)^{3/2}}{2Wv(1+v)^2}.$$

Thus, if λ satisfies the condition above, then expected number of demands in the environment is finite and the TSP-based policy is stable.

We now compute an upper bound on the steady state expected time a demand spends in the environment. If we denote $a := \frac{v}{1+v} + \sqrt{\frac{2Wv\lambda}{(1-v^2)^{3/2}}}$, $b := \frac{5W}{2\sqrt{1-v^2}}$ and $\bar{Y} := \lim_{i\to+\infty} \mathbb{E}[Y_i]$, then the recurrence Eq. (5) implies

$$\bar{Y} \leq \frac{vb}{1-a}.$$

Thus, in the steady state, the vehicle would be at a distance of at most $vb/(1-a)$ from the generator in expected value. Suppose the jth demand arrived between time iterations $i-1$ and i, i.e., in the time interval $(t_{i-1}, t_i]$. Then, the distance traveled by demand i before being serviced is at most Y_{i+1}, assuming that it is the first demand to arrive in the interval $(t_{i-1}, t_i]$ and the last among them to be serviced. Thus, the time spent by the jth demand in the environment (i.e., the delay) satisfies

$$D_j \leq Y_{i+1}/v.$$

Taking expectation, as $j \to +\infty$, we must also have $i \to +\infty$; otherwise, if $i \to i^* < +\infty$, then it would mean that there are infinite number of demands which arrive in $(t_{i^*-1}, t_{i^*}]$, contradicting the fact that the system is stable. Thus, the steady state expected time spent by a demand in the environment satisfies

$$\lim_{j\to+\infty} \mathbb{E}[D_j] \leq \lim_{i\to+\infty} \mathbb{E}[Y_{i+1}]/v = \bar{Y}/v \leq \frac{b}{1-a},$$

which is bounded if $\lambda < (1-v^2)^{3/2}/2Wv(1+v)^2$. ■

B. Limiting Case of Low Speed Demands

In this section we focus on the case when $\lambda \to +\infty$ and, by the necessary stability condition in Theorem IV.1, $v \to 0^+$. Recall that for this case, the sufficient stability condition for the TSP-based policy is that $\lambda < 1/(2vW)$. This differs by a factor of 8 from the policy independent necessary stability condition of $\lambda < 4/(vW)$. By utilizing the tight asymptotic expression for the length of the TSP path, given in Theorem II.2, in place of the bound in Lemma II.1, we can reduce this factor to approximately 2.

To begin, consider an iteration i of the TSP-based policy, and let $Y_i > 0$ be the position of the service vehicle. Then, in the limit as $\lambda \to +\infty$, the number of outstanding demands in that iteration $n_i \to +\infty$. Thus, applying Theorem II.2 and Lemma II.1, the position of the vehicle at the end of the iteration is given by

$$Y_{i+1} = v\beta_{\mathrm{TSP}}\sqrt{n_i A} = v\beta_{\mathrm{TSP}}\sqrt{n_i Y_i W},$$

where $A := Y_i W$ is the area of the region below the vehicle at the ith iteration. Thus, conditioned on Y_i being bounded away from 0, we have

$$\begin{aligned}\mathbb{E}[Y_{i+1}|Y_i] &= v\beta_{\mathrm{TSP}}\mathbb{E}\left[\sqrt{Wn_iY_i}\right]\\ &\leq v\beta_{\mathrm{TSP}}\sqrt{WY_i\mathbb{E}[n_i]},\end{aligned}$$

where we have applied Jensen's inequality. Using Lemma IV.3, $\mathbb{E}[n_i] = WY_i\lambda/(vW)$ and thus

$$\mathbb{E}[Y_{i+1}|Y_i] \leq v\beta_{\mathrm{TSP}}\sqrt{W^2Y_i^2\frac{\lambda}{vW}} = \beta_{\mathrm{TSP}}\sqrt{\lambda vW}Y_i.$$

Thus, the sufficient condition for stability of the TSP-based policy as $\lambda \to +\infty$ (and thus $v \to 0^+$) is

$$\lambda < \frac{1}{\beta_{\mathrm{TSP}}^2 vW} \approx \frac{2}{vW},$$

where $\beta_{\mathrm{TSP}} \approx 0.712$. Hence, in the limiting regime as $\lambda \to +\infty$, the sufficient stability condition for the TSP-based policy is within a constant factor of the policy independent necessary condition.

VI. Simulations

In this section we present results of numerical experiments of the TSP-based policy. The `linkern`[1] solver was used to generate approximations to the optimal TSP tour. We use these experiments to study how the steady-state expected delay varies with the generation rate of the demands for different speeds. The policy was simulated for a fixed number of iterations, large enough to ensure that steady-state was reached. We obtained the steady-state delay by computing the mean of the delay in the last 20 iterations. This was repeated 10 times to obtain an estimate of the steady-state expected delay for one value of the generation rate and for a given speed. We then repeat the experiment for three different values of the speed. The variation of the expected delay with the generation rate is presented in Fig. 5 for different values of speed.

We observe that for sufficiently large values of v such as $v = 0.2$, the necessary condition on the arrival rate is almost 4 times that of the arrival rate that causes instability. However, this ratio decreases as v reduces. Specifically, for $v = 0.05$, the ratio is approximately 2.5. These observations are consistent with our theoretical analysis that predicts that in the limit as $v \to 0^+$, the ratio of the necessary to sufficient condition on λ tends to 2.

VII. Conclusions and Future directions

We introduced a vehicle routing problem in which a service vehicle seeks to serve demands that arrive via a Poisson process on a line segment and that move with a fixed speed in a direction perpendicular to the line. For the existence of a stabilizing service policy, we first derived a necessary condition on the arrival rate of the demands as a function of the speed ratio between the demands and the vehicle, and

[1] `linkern` is written in ANSI C and is freely available for academic research use at `http://www.tsp.gatech.edu//concorde.html`.

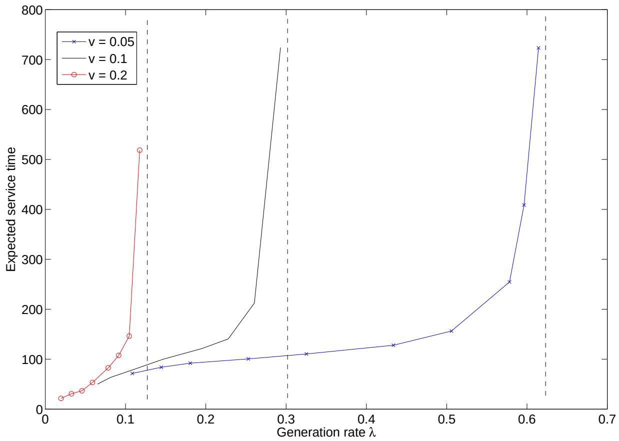

Fig. 5. Simulation results for TSP-based policy: variation of the steady-state expected delay with the arrival rate for three different values of the demand speeds. A dotted line is an asymptote to a curve, i.e., the extrapolated value of λ that leads to instability for a given v. The length of the generator is $W = 50$.

the length of the line segment. Then, we proposed a novel service policy for the vehicle which involves sequentially servicing all the outstanding demands as per the translational TSP through the moving demands. We derived a sufficient condition on the arrival rate of the demands for stability of the TSP-based policy. In the limiting case of the relative speed tending to zero, we showed that the necessary and the sufficient conditions on the arrival rate are within a constant factor. We also provided an upper bound on the expected time spent by the demand in the environment before being served. In companion paper [10], we analyze the first-come-first-served (FCFS) policy and show that in the regimes of high demand speeds, the policy is stable for higher arrival rates than the TSP-based policy. Further, we show that in the high demand speed regime, the FCFS is the optimal policy.

In the future, we envision to address versions of the present problem involving multiple service vehicles. Another interesting direction is to consider non-uniform spatial arrival of the demands on the generating line.

Acknowledgments

This material is based upon work supported in part by ARO-MURI Award W911NF-05-1-0219, ONR Award N00014-07-1-0721 and by the Institute for Collaborative Biotechnologies through the grant DAAD19-03-D-0004 from the U.S. Army Research Office.

References

[1] D. J. Bertsimas and G. J. van Ryzin, "A stochastic and dynamic vehicle routing problem in the Euclidean plane," *Operations Research*, vol. 39, pp. 601–615, 1991.

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

[3] J. D. Papastavrou, "A stochastic and dynamic routing policy using branching processes with state depended immigration," *European Journal of Operational Research*, vol. 95, pp. 167–177, 1996.

[4] E. Frazzoli and F. Bullo, "Decentralized algorithms for vehicle routing in a stochastic time-varying environment," in *IEEE Conf. on Decision and Control*, Paradise Island, Bahamas, Dec. 2004, pp. 3357–3363.

[5] B. Korte and J. Vygen, *Combinatorial Optimization: Theory and Algorithms*, 3rd ed., ser. Algorithmics and Combinatorics. New York: Springer Verlag, 2005, no. 21.

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

[7] Y. Asahiro, E. Miyano, and S. Shimoirisa, "Grasp and delivery for moving objects on broken lines," *Theory of Computing Systems*, vol. 42, no. 3, pp. 289–305, 2008.

[8] M. Hammar and B. J. Nilsson, "Approximation results for kinetic variants of TSP," *Discrete and Computational Geometry*, vol. 27, no. 4, pp. 635–651, 2002.

[9] C. S. Helvig, G. Robins, and A. Zelikovsky, "The moving-target traveling salesman problem," *Journal of Algorithms*, vol. 49, no. 1, pp. 153–174, 2003.

[10] S. L. Smith, S. D. Bopardikar, F. Bullo, and J. P. Hespanha, "Dynamic vehicle routing with moving demands – Part II: High speed demands or low arrival rates," in *American Control Conference*, St. Louis, MO, June 2009, submitted.

[11] L. Few, "The shortest path and the shortest road through n points in a region," *Mathematika*, vol. 2, pp. 141–144, 1955.

[12] J. Beardwood, J. Halton, and J. Hammersly, "The shortest path through many points," in *Proceedings of the Cambridge Philosophy Society*, vol. 55, 1959, pp. 299–327.

[13] G. Percus and O. C. Martin, "Finite size and dimensional dependence of the Euclidean traveling salesman problem," *Physical Review Letters*, vol. 76, no. 8, pp. 1188–1191, 1996.

[14] R. Isaacs, *Differential games*. New York: John Wiley, 1965.

Appendix

In this Appendix we present the proof of Lemma II.1.

Proof: Suppose the rectangular region is given by $0 \le x \le 1$, $0 \le y \le h$. Let m be a positive integer (to be chosen later) and let the n points be denoted by $\{\mathbf{q}_1, \dots, \mathbf{q}_n\}$. We now construct two paths through the points. The first consists of (a) the $m+1$ lines $y = 0, h/m, 2h/m, \dots, h$; (b) the n shortest distances from each of the n points to the nearest such line, each traveled twice, and (c) suitable portions of the lines $x = 0$, $0 \le y \le h$, and $x = 1$, $0 \le y \le h$. This is illustrated in Fig. 6. The length of this path is

$$l_1 = m + 1 + 2\sum_{i=1}^{n} d_1(\mathbf{q}_i) + h,$$

where the notation $d_1(\mathbf{q}_i)$ denotes the shortest distance of point $\mathbf{q}_i$ from the nearest of the $m+1$ lines. The second path is constructed similarly using the m lines $y = h/2m, 3h/2m, \dots, (2m-1)h/2m$. This path also commences on $y = h$, passes through the above m lines (visiting the points whenever they are at the shortest distance from these m lines) and ends on $y = 0$. The length of this path is

$$l_2 = (m+2) + 2\sum_{i=1}^{n} d_2(\mathbf{q}_i) + h,$$

where the notation $d_2(\mathbf{q}_i)$ denotes the shortest distance of point $\mathbf{q}_i$ from the nearest of the new m lines.

Observe that $d_1(\mathbf{q}_i) + d_2(\mathbf{q}_i) = h/2m$. Hence,

$$l_1 + l_2 = 2m + 3 + 2h + hn/m.$$

Now choose m to be the integer nearest to $\sqrt{hn/2}$, so that

$n = 2(m+\theta)^2/h$, where $|\theta| \leq 1$. Thus,

$$\begin{aligned} l_1 + l_2 &= 2m + 3 + 2h + 2(m+\theta)^2/m \\ &= 4(m+\theta) + 2h + 3 + 2\theta^2/m \\ &\leq 2\sqrt{2hn} + 2h + 5. \end{aligned}$$

Thus, at least one of the two paths must have length upper bounded by $\sqrt{2hn} + h + 5/2$. ■

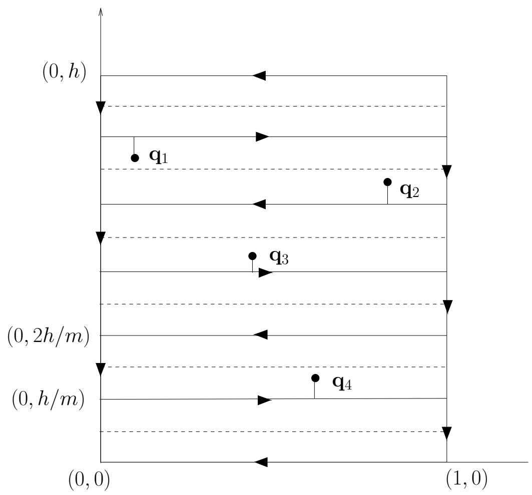

Fig. 6. Illustration of the proof of Theorem II.1. The dots indicate the locations of the points inside a rectangle of size $1 \times h$. The first of the two paths considered in the proof through the points begins at $(1, h)$ and follows the direction of the arrows, visiting a point whenever it is within a distance of $h/2m$ for a specific integer m from the solid horizontal lines.