

A Dynamic Boundary Guarding Problem with Translating Targets

Stephen L. Smith Shaunak D. Bopardikar Francesco Bullo

Abstract—We introduce a problem in which a service vehicle seeks to defend a deadline (boundary) from dynamically arriving mobile targets. The environment is a rectangle and the deadline is one of its edges. Targets arrive continuously over time on the edge opposite the deadline, and move towards the deadline at a fixed speed. The goal for the vehicle is to maximize the fraction of targets that are captured before reaching the deadline. We consider two cases; when the service vehicle is faster than the targets, and; when the service vehicle is slower than the targets. In the first case we develop a novel vehicle policy based on computing longest paths in a directed acyclic graph. We give a lower bound on the capture fraction of the policy and show that the policy is optimal when the distance between the target arrival edge and deadline becomes very large. We present numerical results which suggest near optimal performance away from this limiting regime. In the second case, when the targets are slower than the vehicle, we propose a policy based on servicing fractions of the translational minimum Hamiltonian path. In the limit of low target speed and high arrival rate, the capture fraction of this policy is within a small constant factor of the optimal.

I. INTRODUCTION

Vehicle motion planning in dynamic environments arises in many important autonomous vehicle applications. In areas such as environmental monitoring, surveillance and perimeter defence, the vehicle must re-plan its motion as it acquires information on its surroundings. In addition, remote operators may add tasks to, or remove tasks from, the vehicle’s mission in real-time. In this paper we consider a problem in which a vehicle must defend a boundary in a dynamic environment with approaching targets.

Static vehicle routing problems consider planning a path through a fixed number of locations. Examples include the traveling salesperson problem (TSP) [1], the deadline-TSP and vehicle routing with time-windows [2]. Recently, researchers have looked at the TSP with moving objects. In [3] the authors consider objects moving on straight lines and focus on the case when the objects are slower than the vehicle and when the vehicle moves parallel to the x - or y -axis. The same problem is studied in [4], but with arbitrary vehicle motion, and it is called the translational TSP. The authors of [4] propose a polynomial-time approximation scheme to catch all objects in minimum time. Other variations of the problem are studied in [5] and [6].

Dynamic vehicle routing (DVR) is a class of problems in which vehicles must plan paths through service demand locations that arrive sequentially over time. An early DVR

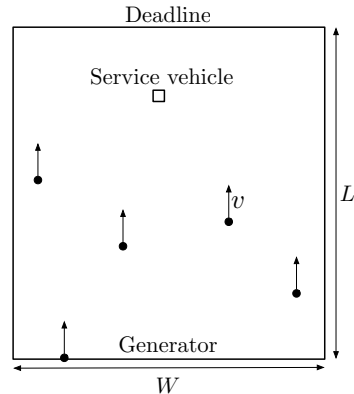


Fig. 1. The problem setup. Demands are shown as black disks approaching the deadline at speed v . The service vehicle is a square.

problem was the dynamic traveling repairperson problem [7], [8], where each demand assumes a fixed location upon arrival, and the vehicle must spend some amount of on-site service time at each location. This problem has also been studied from the online algorithm perspective [9], [10]. Other recent DVR problems include DVR with demands that disappear if left unserved for a certain amount of time [11], and demands with different priority levels [12]. In our earlier work [13], we introduced a DVR problem in which demands arrive on a line segment and move in a perpendicular direction at a fixed speed slower than the vehicle. We derived conditions on the demand arrival rate and demand speed for the existence of a vehicle routing policy which can serve all demands, and a proposed a policy based on the translational minimum Hamiltonian path.

Contributions: In this paper we introduce the following problem (see Fig. 1): Targets (or demands) arrive according to a stochastic process on a line segment of length W . Upon arrival the demands move with fixed speed v towards a deadline which is at a distance L from the generator. A unit speed service vehicle seeks to capture the demands before they reach the deadline (i.e., within L/v time units of being generated). The performance metric is the fraction of demands that are captured before reaching the deadline.

We assume that the arrival process is uniform along the line segment and temporally Poisson with rate λ . In the case when the demands are faster than the service vehicle (i.e., $v \geq 1$) we introduce the novel Longest Path policy, which is based on computing longest paths in a directed acyclic *reachability graph*. When $L \geq vW$, we derive a lower bound on the capture fraction as a function of the system parameters. We show that the Longest Path policy is the optimal policy when L is much greater than vW . In the

This material is based upon work supported in part by ARO-MURI Award W911NF-05-1-0219 and ONR Award N00014-07-1-0721.

S. L. Smith, S. D. Bopardikar, and F. Bullo are all with the Center for Control, Dynamical Systems and Computation, University of California at Santa Barbara, Santa Barbara, CA 93106, USA. Email: {stephen, shaunak, bullo}@engineering.ucsb.edu.

case when the demands are slower than the service vehicle (i.e., $v < 1$), we propose a policy based on the translational minimum Hamiltonian path called the TMHP-fraction policy. In the limit of low demand speed and high arrival rate, the capture fraction of this policy is within a small constant factor of the optimal. We present numerical simulations which verify our results, and show that the Longest Path policy performs very near the optimal even when $L < vW$.

The paper is organized as follows. In Section II we formulate the problem and in Section III we review some background material. In Section IV we consider the case of $v \geq 1$ and introduce the Longest Path Policy. In Section V we study $v < 1$ and introduce the TMHP-fraction policy. Finally, in Section VI we present simulation results. Due to page constraints we omit the proofs of all results. An extended version of this paper, including all proofs, can be found in [14].

II. PROBLEM FORMULATION

Consider an environment $\mathcal{E} := [0, W] \times [0, L] \subset \mathbb{R}^2$ as shown in Figure 1. The line segment $[0, W] \times \{0\} \subset \mathcal{E}$ is termed the *generator*, and the segment $[0, W] \times \{L\} \subset \mathcal{E}$ is termed the *deadline*. The environment contains a single vehicle with position $\mathbf{p}(t) = [X(t), Y(t)]^T \in \mathcal{E}$, modeled as a first-order integrator with unit speed. Demands (or targets) arrive in the environment according to a temporal Poisson process with rate $\lambda > 0$. Upon arrival, each demand assumes a uniformly distributed location on the generator, and then moves with constant speed $v > 0$ in the positive y -direction towards the deadline. If the vehicle intercepts a demand before the demand reaches the deadline, then the demand is captured. On the other hand, if the demand reaches the deadline before being intercepted by the vehicle, then the demand escapes. Thus, to capture a demand, it must be intercepted within L/v time units of being generated.

We let $\mathcal{Q}(t) \subset \mathcal{E}$ denote the set of all outstanding demand locations at time t . If the i th demand to arrive is captured, then it is removed from \mathcal{Q} and placed in the set $\mathcal{Q}_{\text{capt}}$ with cardinality n_{capt} . If the i th demand escapes, then it is removed from \mathcal{Q} and placed in \mathcal{Q}_{esc} with cardinality n_{esc} .

Causal Policy: A causal feedback control policy for the vehicle is a map $\mathcal{P} : \mathcal{E} \times \mathbb{F}(\mathcal{E}) \rightarrow \mathbb{R}^2$, where $\mathbb{F}(\mathcal{E})$ is the set of finite subsets of \mathcal{E} , assigning a commanded velocity to the service vehicle as a function of the current state of the system: $\dot{\mathbf{p}}(t) = \mathcal{P}(\mathbf{p}(t), \mathcal{Q}(t))$.

Non-causal Policy: In a non-causal feedback control policy the commanded velocity of the service vehicle is a function of the current and future state of the system. Such policies are not physically realizable, but they will prove useful in the upcoming analysis.

Formally, let the generation of demands commence at time $t = 0$, and consider the sequence of demands (q_1, q_2, \dots) arriving at increasing times (t_1, t_2, \dots) , with x -coordinates (x_1, x_2, \dots) . We can also model the arrival process by assuming that at time $t = 0$, all demands are located in $[0, W] \times (-\infty, 0]$, move in the y -direction at speed v for all $t > 0$, and are revealed to the service vehicle when they cross the generator. Thus, at time $t = 0$, the position of the

i th demand is $(x_i, v(t - t_i))$. We can define a set containing the position of all demands in the region $[0, W] \times (-\infty, 0]$ at time t as $\mathcal{Q}_{\text{unarrived}}(t)$. Then, a non-causal policy is one for which $\dot{\mathbf{p}}(t) = \mathcal{P}(\mathbf{p}(t), \mathcal{Q}(t) \cup \mathcal{Q}_{\text{unarrived}}(t))$.

Problem Statement: The goal in this paper is to find causal policies P that maximize the fraction of demands that are captured $\mathbb{F}_{\text{cap}}(P)$, termed the *capture fraction*, where

$$\mathbb{F}_{\text{cap}}(P) := \limsup_{t \rightarrow +\infty} \mathbb{E} \left[\frac{n_{\text{capt}}(t)}{n_{\text{capt}}(t) + n_{\text{esc}}(t)} \right].$$

III. PRELIMINARY COMBINATORIAL RESULTS

We now review the longest path problem and optimal tours/paths through a set of points.

A. 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 vertex 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 [15]. However, if the graph is a DAG, then the longest path problem has an efficient dynamic programming solution [1] with complexity $O(|V| + |E|)$, that relies on topologically sorting [16] the vertices.

B. The Euclidean Shortest Path/Tour Problems

Given a set \mathcal{Q} of n points in \mathbb{R}^2 , the Euclidean traveling salesperson problem (ETSP) is to find the minimum-length tour (i.e., cycle) of \mathcal{Q} . Letting $\text{ETSP}(\mathcal{Q})$ denote the minimum length of a tour of \mathcal{Q} , we can state the following result.

Theorem III.1 (Length of ETSP tour, [17]) *Consider a set \mathcal{Q} of n points independently and uniformly distributed in a compact set \mathcal{E} of area $|\mathcal{E}|$. Then, there exists a constant β_{TSP} such that, with probability one,*

$$\lim_{n \rightarrow +\infty} \frac{\text{ETSP}(\mathcal{Q})}{\sqrt{n}} = \beta_{\text{TSP}} \sqrt{|\mathcal{E}|}. \quad (1)$$

The constant β_{TSP} has been estimated numerically as $\beta_{\text{TSP}} \approx 0.7120 \pm 0.0002$, [18].

The Euclidean Minimum Hamiltonian Path (EMHP) problem is to compute the shortest path through a set of points. In this paper we consider a constrained EMHP problem: Given a start point \mathbf{s} , a set of n points \mathcal{Q} , and a finish point \mathbf{f} , all in \mathbb{R}^2 , determine the shortest path which starts at \mathbf{s} , visits each point in \mathcal{Q} exactly once, and terminates at \mathbf{f} . We let $\text{EMHP}(\mathbf{s}, \mathcal{Q}, \mathbf{f})$ denote the length of the shortest path.

Corollary III.2 (Length of EMHP) *Consider a set \mathcal{Q} of n points independently and uniformly distributed in a compact*

set \mathcal{E} of area $|\mathcal{E}|$, and any two points $\mathbf{s}, \mathbf{f} \in \mathcal{E}$. Then with probability one,

$$\lim_{n \rightarrow +\infty} \frac{\text{EMHP}(\mathbf{s}, \mathcal{Q}, \mathbf{f})}{\sqrt{n}} = \beta_{\text{TSP}} \sqrt{|\mathcal{E}|},$$

where β_{TSP} is defined in Theorem III.1.

The above corollary states that the length of the EMHP and the ETSP tour are asymptotically equal, and it follows directly from the fact that as $n \rightarrow +\infty$, the diameter of \mathcal{E} is negligible when compared to the length of the tour/path.

C. Translational Minimum Hamiltonian Path (TMHP)

The TMHP problem is posed as follows. Given initial coordinates; \mathbf{s} of a start point, $\mathcal{Q} := \{\mathbf{q}_1, \dots, \mathbf{q}_n\}$ of a set of points, and \mathbf{f} of a finish point, all moving with speed $v \in]0, 1[$ in the positive y -direction, determine a minimum length path that starts at time zero from point \mathbf{s} , visits all points in the set \mathcal{Q} and ends at the finish point. The following gives a solution [4] for the TMHP problem.

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

$$g(x, y) = \left(\frac{x}{\sqrt{1-v^2}}, \frac{y}{1-v^2} \right).$$

(ii) Compute the EMHP that starts at $g_v(\mathbf{s})$, passes through $\{g(\mathbf{q}_1), \dots, g(\mathbf{q}_n)\} =: g(\mathcal{Q})$ and ends at $g(\mathbf{f})$.

(iii) To reach a translating point with initial position (x, y) from the initial position (X, Y) , move towards the point $(x, y + vT)$, where

$$T = \frac{\sqrt{(1-v^2)(X-x)^2 + (Y-y)^2}}{1-v^2} - \frac{v(Y-y)}{1-v^2}.$$

The length $\text{TMHP}_v(\mathbf{s}, \mathcal{Q}, \mathbf{f})$ of the path is as follows.

Lemma III.3 (TMHP length, [4]) *Let the initial coordinates $\mathbf{s} = (x_s, y_s)$ and $\mathbf{f} = (x_f, y_f)$, and the speed of the points $v \in]0, 1[$. Then,*

$$\text{TMHP}_v(\mathbf{s}, \mathcal{Q}, \mathbf{f}) = \text{EMHP}(g(\mathbf{s}), g(\mathcal{Q}), g(\mathbf{f})) + \frac{v(y_f - y_s)}{1-v^2}.$$

IV. DEMAND SPEED GREATER THAN VEHICLE SPEED

Here we develop a policy for the case when the demand speed $v \geq 1$. In this policy, the service vehicle remains on the deadline and services demands as per the longest path in a directed acyclic reachability graph. In this section we begin by introducing the reachability graph, and then proceed to state and analyze the Longest Path policy.

A. Reachable Demands

Consider a demand generated at time $t_1 \geq 0$ at position $(x, 0)$. The demand moves in the positive y -direction at speed $v \geq 1$, and thus $(x(t), y(t)) = (x, v(t - t_1))$ for each $t \in [t_1, T]$, where T is either the time of escape (i.e., $T = L/v + t_1$), or it is the time of capture. Now, given the service vehicle location $(X(t), Y(t))$, a demand with position $(x, y(t))$ is reachable if and only if

$$v|X(t) - x| \leq Y(t) - y(t). \quad (2)$$




Fig. 2. The construction of the reachability graph. The top-left figure shows the set of reachable points from a vehicle positioned on the deadline. The top-right and bottom-left figures show the reachable set from demand locations. The bottom-right figure shows the reachability graph.

That is, the service vehicle must be at a height of at least $v|X(t) - x|$ above the demand in order to capture it.

Definition IV.1 (Reachable set) *The reachable set from a position $(X, Y) \in \mathcal{E}$ is*

$$R(X, Y) := \{(x, y) \in \mathcal{E} : v|X - x| \leq |Y - y|\}.$$

If the service vehicle is located at (X, Y) , then a demand can be captured if and only if it lies in the set $R(X, Y)$.

An example of the reachable set is shown in Figure 2. Next, given a demand in the reachable set, the following motion gives a method of capture.

Definition IV.2 (Intercept motion) *Consider a vehicle position $((X(\bar{t}), Y(\bar{t}))$ and a demand position $(x, y(\bar{t})) \in R(X(\bar{t}), Y(\bar{t}))$ at time $\bar{t} \geq 0$. In intercept motion, the service vehicle captures the demand by first moving horizontally at unit speed to the position $(x_i, Y(\bar{t}))$, and then waiting at the location for the demands arrival.*

It turns out that intercept motion is optimal, as summarized in the following lemma. The proof is contained in [14].

Lemma IV.3 (Optimality of intercept motion) *Consider $v \geq 1$, and let the service vehicle be initially positioned on the deadline. Then, there is an optimal policy in which the service vehicle uses only intercept motion.*

Next, consider the set of demands in $R(X(\bar{t}), Y(\bar{t}))$, and suppose the vehicle chooses to capture demand i , with position $\mathbf{q}_i(\bar{t}) = (x_i, y_i(\bar{t})) \in R(X(\bar{t}), Y(\bar{t}))$. Upon capture

The Greedy Path (GP) policy

Assumes: Vehicle is located at (X, L)

- 1 Compute the reachability set $R(X, L)$.
 - 2 Capture the demand in $R(X, L)$ with the highest y -coordinate using intercept motion.
 - 3 Repeat.
-

Given a set of outstanding demands $\mathcal{Q}(t)$ at time t , the Greedy Path policy generates a suboptimal longest path through $\mathcal{Q}(t)$. In addition, the vehicle position is independent of all outstanding demands, except the demand currently being captured. Thus, the capture fraction of the Greedy Path policy provides a lower bound for the capture fraction of the Longest Path policy. We are now able to establish the following result, whose proof is contained in [14].

Theorem IV.8 (Lower Bound for Longest Path policy)

If $L \geq vW$, then for the Longest Path policy

$$\mathbb{F}_{\text{cap}}(\text{LP}) \geq \mathbb{F}_{\text{cap}}(\text{GP}) \geq \frac{1}{\sqrt{\pi\alpha} \operatorname{erf}(\sqrt{\alpha}) + e^{-\alpha}},$$

where $\alpha = \lambda W/2$ and $\operatorname{erf} : \mathbb{R} \rightarrow [-1, 1]$ is the error function.

V. DEMAND SPEED LESS THAN VEHICLE SPEED

In this section we study the case when the demand speed $v < 1$. For this case, an upper bound on the capture fraction has been derived in [13]. We introduce a policy which is a variant of the TMHP-based policy in [13], and lower bound its capture fraction in the limit of low demand speed and high demand arrival rate.

A. Capture Fraction Upper Bound

The following theorem upper bounds the capture fraction of every policy for the case of $v < 1$.

Theorem V.1 (Capture fraction upper bound, [13]) If $v < 1$, then for every causal policy P

$$\mathbb{F}_{\text{cap}}(P) \leq \min \left\{ 1, \frac{2}{\sqrt{v\lambda W}} \right\}.$$

The proof of the above theorem is contained in [13], and relies on a computation of the expected minimum distance between demands. Notice that for low demands speed, i.e., $v \ll 1$, it may be possible to achieve a capture fraction of one, even for high arrival rates.

B. The TMHP-fraction Policy

In Section III-C we reviewed the translational minimum Hamiltonian Path (TMHP) through a set of demands. The following policy utilizes this path to service demands. Figure 4 shows an example of the TMHP-fraction policy. In contrast with the LP policy, where the vehicle remains on the deadline, in the TMHP-fraction policy the vehicle follows the TMHP using minimum time motion between demands as described in Section III-C.

Notice that none of the demands in the region $[0, W] \times [0, L/2]$ at time t will have escaped before time $t + L/(2v)$.




Fig. 4. The TMHP-fraction policy. The left-hand figure shows a TMHP through all outstanding demands. The right-figure shows the instant when the vehicle has followed the path for $L/(2v)$ time units and recomputes its path, allowing some demands to escape.

Thus, the vehicle is guaranteed that for the first $L/(2v)$ time units, all demands on the TMHP path are still in the environment. For the TMHP-fraction policy we have the following result (see [14] for its proof).

The TMHP-fraction (TF) policy

Assumes: Vehicle is located on the line $y = L/2$.

- 1 Compute a translational minimum Hamiltonian path through all outstanding demands in $[0, W] \times [0, L/2]$, starting at the service vehicle position, and terminating at the demand with the lowest y -coordinate.
 - 2 **if** time to travel entire path is less than $L/(2v)$ **then**
 - 3 Service all outstanding demands by following the computed path.
 - 4 **else**
 - 5 Service outstanding demands along the computed path for $L/(2v)$ time units.
 - 6 Repeat.
-

Theorem V.2 (TMHP-fraction policy lower bound) In the limit as $v \rightarrow 0^+$ and $\lambda \rightarrow +\infty$, the capture fraction of the TMHP-fraction policy satisfies

$$\mathbb{F}_{\text{cap}}(\text{TF}) \geq \min \left\{ 1, \frac{1}{\beta_{\text{TSP}} \sqrt{v\lambda W}} \right\}.$$

Remark V.3 (Bound comparison) In the limit as $v \rightarrow 0^+$, and $\lambda \rightarrow +\infty$, the capture fraction of the TMHP-fraction policy is within a factor of $2\beta_{\text{TSP}} \approx 1.42$ of the optimal. \square

VI. SIMULATIONS

We now present two sets of results from numerical experiments. The first set compares the Longest Path policy with $\eta = 1$ to the Non-causal Longest Path policy and to the theoretical lower bound in Theorem IV.8. The second set compares the TMHP-fraction policy to the policy independent upper bound in Theorem V.1 and the lower bound in Theorem V.2.

To simulate the LP and the NCLP policies, we perform 10 runs of the policy, where each run consists of 5000 demands. A comparison of the capture fractions for the two policies is presented in Figure 5. When $L > vW$, the capture fraction of the LP policy is nearly identical to that of the NCLP policy.




Fig. 5. Simulation results for LP policy (solid red line with error bars showing \pm one standard deviation) and the NCLP policy (dashed black line) for an environment of width $W = 120$ and length $L = 500$. In (a), $L > vW$, and the lower bound in Theorem IV.8 is shown in solid green.




Fig. 6. Simulation results for TMHP-fraction policy. The solid black curve shows the upper bound in Theorem V.1 and the dashed line shows the lower bound in Theorem V.2. Numerical results are shown with error bars.

Even in Figure 5(a), where $L < vW$, the capture fraction of the LP policy is within 2% of the NCLP policy, and thus the optimal. This suggests that the Longest Path policy is essentially optimal over a large range of parameter values.

To simulate the TMHP-fraction policy, the `linkern`¹ solver is used to generate approximations to the optimal TMHP. For each value of arrival rate, we determine the capture fraction by taking the mean over 10 runs of the policy. A comparison of the simulation results with the theoretical results from Section V are presented in Figure 6. For $v = 0.01$ in Fig. 6(a), the experimental results are in near exact agreement with the theoretical lower bound in Theorem V.1. For $v = 0.05$ in Fig. 6(b), the experimental results are within 5% of the theoretical lower bound. However, notice that the experimental capture fraction is smaller than the theoretical lower bound. This is due to several factors. First, we have not reached the limit as $v \rightarrow 0^+$ and $\lambda \rightarrow +\infty$ where the asymptotic value of $\beta_{\text{TSP}} \approx 0.712$ holds. Second, we are using an approximate solution to the optimal TMHP, generated via the `linkern` solver.

VII. CONCLUSIONS

In this paper we introduced a pursuit problem in which a vehicle must defend a deadline from approaching demands.

¹`linkern` is freely available for academic research use at <http://www.tsp.gatech.edu/concorde.html>.

We presented novel policies in the case when the demand speed is greater than the vehicle speed, and in the case when the demand speed is less than the vehicle speed. In the former case we introduced the Longest Path policy which is based on computing longest paths in the directed acyclic reachability graph, and in the latter case we introduced the TMHP-fraction policy. For each policy, we analyzed the fraction of demands that are captured.

There are many areas for future work. The Longest Path policy has promising extensions to the case when demands have different priority levels, and to the case of multiple vehicles. We would also like to fully characterize the capture fraction when $L < vW$, and tighten our existing bounds to reflect the near optimal performance shown in simulation.

REFERENCES

- [1] N. Christofides, *Graph Theory: An Algorithmic Approach*. Academic Press, 1975.
- [2] N. Bansal, A. Blum, S. Chawla, and A. Meyerson, "Approximation algorithms for deadline-TSP and vehicle routing with time-windows," in *ACM Symposium on the Theory of Computing*, Chicago, IL, 2004, pp. 166–174.
- [3] P. Chalasani and R. Motwani, "Approximating capacitated routing and delivery problems," *SIAM Journal on Computing*, vol. 28, no. 6, pp. 2133–2149, 1999.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] H. N. Psaraftis, "Dynamic vehicle routing problems," in *Vehicle Routing: Methods and Studies*, B. Golden and A. Assad, Eds. Elsevier (North-Holland), 1988, pp. 223–248.
- [8] 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.
- [9] A. Borodin and R. El-Yaniv, *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
- [10] S. O. Krumke, W. E. de Paepe, D. Poensgen, and L. Stougie, "News from the online traveling repairman," *Theoretical Computer Science*, vol. 295, no. 1-3, pp. 279–294, 2003.
- [11] 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*, 2008, to appear. [Online]. Available: <http://ares.lids.mit.edu/papers/Pavone.Bisnik.ea.MONE08.pdf>
- [12] 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*, Feb. 2009, submitted.
- [13] 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*, Aug. 2009. (Submitted, Mar 2009) to appear.
- [14] S. L. Smith, S. D. Bopardikar, and F. Bullo, "A dynamic boundary guarding problem with translating targets," Aug. 2009, available at <http://arxiv.org/abs/0908.3929>.
- [15] B. Korte and J. Vygen, *Combinatorial Optimization: Theory and Algorithms*, 4th ed., ser. Algorithmics and Combinatorics. Springer, 2007, vol. 21.
- [16] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. MIT Press, 2001.
- [17] 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.
- [18] A. 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.