

Dynamic Multi-Vehicle Routing with Multiple Classes of Demands

Marco Pavone Stephen L. Smith Francesco Bullo Emilio Frazzoli

Abstract—In this paper we study a dynamic vehicle routing problem in which there are multiple vehicles and multiple classes of demands. Demands of each class arrive in the environment randomly over time and require a random amount of on-site service that is characteristic of the class. To service a demand, one of the vehicles must travel to the demand location and remain there for the required on-site service time. The quality of service provided to each class is given by the expected delay between the arrival of a demand in the class, and that demand’s service completion. The goal is to design a routing policy for the service vehicles which minimizes a convex combination of the delays for each class. First, we provide a lower bound on the achievable values of the convex combination of delays. Then, we propose a novel routing policy and analyze its performance under heavy load conditions (i.e., when the fraction of time the service vehicles spend performing on-site service approaches one). The policy performs within a constant factor of the lower bound (and thus the optimal), where the constant depends only on the number of classes, and is independent of the number of vehicles, the arrival rates of demands, the on-site service times, and the convex combination coefficients.

I. INTRODUCTION

Consider a bounded environment \mathcal{E} in the plane which contains n service vehicles. Demands for service arrive in \mathcal{E} sequentially over time and each demand is a member of one of m classes. Upon arrival, a demand assumes a location in \mathcal{E} , and requires a class dependent amount of on-site service time. To service a demand, one of the n vehicles must travel to the demand location and perform the on-site service. If we specify a policy by which the vehicles serve demands, then the expected delay for demands of class α , denoted D_α , is the expected amount of time between a demands arrival and its service completion. Then, given coefficients $c_1, \dots, c_m > 0$, the goal is to find the vehicle routing policy that minimizes

$$c_1 D_1 + \dots + c_m D_m.$$

By increasing the coefficients for certain classes, a higher priority level can be given to their demands. This problem, which we call *dynamic vehicle routing with priority classes*,

This research was partially supported by the National Science Foundation, through grants #0705451 and #0705453, by the Office of Naval Research through grant #N00014-07-1-0721, and by the Air Force Office of Scientific Research through grant #FA9550-07-1-0528. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the supporting organizations.

M. Pavone and E. Frazzoli are with the Laboratory for Information and Decision Systems, Aeronautics and Astronautics Department, Massachusetts Institute of Technology, Cambridge, MA 02139, USA; email: {pavone, frazzoli}@mit.edu. S. L. Smith and F. Bullo are with the Center for Control, Dynamical Systems and Computation, Department of Mechanical Engineering, University of California, Santa Barbara, CA 93106, USA {stephen, bullo}@engineering.ucsb.edu.

has important applications in areas such as UAV surveillance, where targets are given different priority levels based on their urgency or potential importance.

In classical queuing theory (i.e., queuing systems in which the demands are not spatially distributed), the problem of priority queues has received much attention, [1]. In [2] the authors characterize the region of delays that are realizable by a single server. This analysis is performed under the assumption that the customer (demand) interarrival times and service times are distributed exponentially. In [3] the achievable delays are studied in more a general setting known as queuing networks.

If service demands are spatially distributed, then providing service becomes a problem in dynamic vehicle routing (DVR). One of the first DVR problems was the dynamic traveling repairperson problem (DTRP) [4], [5]. The DTRP is the single class version of the dynamic vehicle routing with priority classes problem studied in this paper. In [4], [5], the authors study the expected delay of demands and propose optimal policies in both heavy load (i.e., when the fraction of time the service vehicles spend performing on-site service approaches one), and in light load (i.e., when the fraction of time the service vehicles spends performing on-site service approaches zero). In [6], and [7], decentralized policies are developed for the DTRP. Spatial queuing problems have also been studied in the context of urban operations research [8], where approximations are used to cast the problems in the traditional queuing framework. In our previous paper [9], we introduced and studied dynamic vehicle routing with priority classes for the case of two classes and one vehicle. For this case we derived a lower bound on the achievable delay values and proposed the Randomized Priority policy, which performed within a constant factor of the lower bound, for all convex combination coefficients.

The contributions of this paper are as follows. We extend the dynamic vehicle routing with priority classes problem to n service vehicles and m classes of demands. The extension of our previous analysis to multiple classes of demands is very nontrivial. We derive a new lower bound on the achievable values of the convex combination of delays, and propose a new policy in which each class of demands is served separately from the others. We show that the policy performs with a constant factor of $2m^2$ of the optimal. Thus, the constant factor is independent of the number of vehicles, the arrival rates of demands, the on-site service times, and the convex combination coefficients. We also comment on the source of the gap between the upper and lower bounds.

The paper is organized as follows. In Section II we give some asymptotic properties of the traveling salesperson tour.

In Section II-B we formalize the problem and in Section III we derive a lower bound, and in Section IV we introduce and analyze the Separate Queues policy. Finally, in Section V we present simulation results.

II. BACKGROUND AND PROBLEM STATEMENT

In this section we summarize the asymptotic properties of the Euclidean traveling salesperson tour, and formalize dynamic vehicle routing with priority classes.

A. The Euclidean Traveling Salesperson Problem

Given a set Q of N points in \mathbb{R}^2 , the Euclidean traveling salesperson problem (TSP) is to find the minimum-length tour of Q (i.e., the shortest closed path through all points). Let $\text{TSP}(Q)$ denote the minimum length of a tour through all the points in Q . Assume that the locations of the N points are random variables independently and identically distributed, uniformly in a compact set \mathcal{E} with area $|\mathcal{E}|$; in [10] it is shown that there exists a constant β_{TSP} such that, almost surely,

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

The constant β_{TSP} has been estimated numerically as $\beta_{\text{TSP}} \approx 0.7120 \pm 0.0002$, [11]. The bound in equation (1) holds for all compact sets \mathcal{E} , and the shape of \mathcal{E} only affects the convergence rate to the limit. In [8], the authors note that if \mathcal{E} is “fairly compact [square] and fairly convex”, then equation (1) provides an adequate estimate of the optimal TSP tour length for values of N as low as 15.

B. Problem Statement

Consider a compact environment \mathcal{E} in the plane with area $|\mathcal{E}|$. The environment contains n vehicles, each with maximum speed v . Demands of type $\alpha \in \{1, \dots, m\}$ (also called α -demands) arrive in the environment according to a Poisson process with rate λ_α . Upon arrival, demands assume an independently and uniformly distributed location in \mathcal{E} . An α -demand is serviced when the vehicle spends an on-site service time at the demand location, which is generally distributed with finite mean \bar{s}_α .

Consider the arrival of the i th α -demand. The service delay for the i th demand, $D_\alpha(i)$, is the time elapsed between its arrival and its service completion. The wait time is defined as $W_\alpha(i) := D_\alpha(i) - s_\alpha(i)$, where $s_\alpha(i)$ is the on-site service time required by demand i . A policy for routing the vehicles is said to be *stable* if the expected number of demands in the system for each class is bounded uniformly at all times. A necessary condition for the existence of a stable policy is

$$\rho := \frac{1}{n} \sum_{\alpha=1}^m \lambda_\alpha \bar{s}_\alpha < 1. \quad (2)$$

The *load factor* ρ is a standard quantity in queueing theory [1], and is used to capture the fraction of time the n servers (vehicles) must be busy in any stable policy. In general, it is difficult to study a queueing system for all values of $\rho \in [0, 1)$, and a common technique is to focus on

the limiting regimes of $\rho \rightarrow 1^-$, referred to as the *heavy-load* regime, and $\rho \rightarrow 0^+$, referred to as the *light-load* regime.

Given a stable policy P the steady-state service delay for α -demands is defined as $D_\alpha(P) := \lim_{i \rightarrow +\infty} \mathbb{E}[D_\alpha(i)]$, and the steady-state wait time for α -demands is $W_\alpha(P) := D_\alpha(P) - \bar{s}_\alpha$. Thus, for a stable policy P , the *average delay per demand* is

$$D(P) = \frac{1}{\Lambda} \sum_{\alpha=1}^m \lambda_\alpha D_\alpha(P),$$

where $\Lambda := \sum_{\alpha=1}^m \lambda_\alpha$. The average delay per demand is the standard cost functional for queueing systems with multiple classes of demands. Notice that we can write $D(P) = \sum_{\alpha=1}^m c_\alpha D_\alpha(P)$ with $c_\alpha = \lambda_\alpha / \Lambda$. Thus, we can model priority among classes by allowing any convex combination of D_1, \dots, D_m . If $c_\alpha > \lambda_\alpha / \Lambda$, then the delay of α -demands is being weighted more heavily than in the average case. Thus, the quantity $c_\alpha \Lambda / \lambda_\alpha$ gives the priority of α -demands compared to that given in the average delay case. Without loss of generality we can assume that priority classes are labeled so that

$$\frac{c_1}{\lambda_1} \geq \frac{c_2}{\lambda_2} \geq \dots \geq \frac{c_m}{\lambda_m}, \quad (3)$$

implying that if $\alpha < \beta$ for some $\alpha, \beta \in \{1, \dots, m\}$, then the priority of α -demands is at least as high as that of β -demands. With these definitions, we are now ready to state our problem.

Problem Statement: Let Π be the set of all causal, stable and stationary policies for dynamic vehicle routing with priority classes. Given the coefficients $c_\alpha > 0$, $\alpha \in \{1, \dots, m\}$, with $\sum_{\alpha=1}^m c_\alpha = 1$, and satisfying equation (3), let $D(P) := \sum_{\alpha=1}^m c_\alpha D_\alpha(P)$ be the cost of a policy $P \in \Pi$. Then, the problem is to determine a vehicle routing policy P^* , if one exists, such that

$$D(P^*) = \inf_{P \in \Pi} D(P). \quad (4)$$

We let D^* denote the right-hand side of equation (4). A policy P for which $D(P)/D^*$ is bounded has a *constant-factor guarantee*. If $\limsup_{\rho \rightarrow 1^-} D(P)/D^* = \kappa < +\infty$, then the policy P has a *heavy-load constant-factor guarantee* of κ . In this paper we focus on the heavy-load regime, and look for policies with a heavy-load constant-factor guarantee that is *independent* of the number of vehicles, the arrival rates of demands, the on-site service times, and the convex combination coefficients.

III. LOWER BOUND IN HEAVY LOAD

In this section we present a heavy-load lower bound on the delay in Eq. (4).

Theorem III.1 (Heavy load lower bound) In heavy load ($\rho \rightarrow 1^-$), for every routing policy P ,

$$D(P) \geq \frac{\beta_{\text{TSP}}^2 |\mathcal{E}|}{2n^2 v^2 (1 - \rho)^2} \sum_{\alpha=1}^m \left(c_\alpha + 2 \sum_{j=\alpha+1}^m c_j \right) \lambda_\alpha. \quad (5)$$

where c_1, \dots, c_m satisfy Eq. (3).

Proof: Consider a tagged demand i of type α , and let us quantify its total service requirement. The demand requires on-site service time $s_\alpha(i)$. Let us denote by $d_\alpha(i)$ the distance from the location of the demand served prior to i , to i 's location. In order to compute a lower bound on the wait time, we will allow "remote" servicing of some of the demands. For an α -demand i that can be serviced remotely, the travel distance $d_\alpha(i)$ is zero (i.e., a service vehicle can service the i th α -demand from any location by simply stopping for the on-site service time $s_\alpha(i)$). Thus, the wait time for the modified remote servicing problem provides a lower bound on the wait time for the problem of interest. To formalize this idea, we introduce the variables $r_\alpha \in \{0, 1\}$ for each $\alpha \in \{1, \dots, m\}$. If $r_\alpha = 0$, then α -demands can be serviced remotely. If $r_\alpha = 1$, then α -demands must be serviced on location. We assume that $r_\alpha = 1$ for at least one $\alpha \in \{1, \dots, m\}$. Thus, the total service requirement of α -demand i is $r_\alpha d_\alpha(i) + s_\alpha(i)$. The steady-state expected service requirement is $r_\alpha \bar{d}_\alpha + s_\alpha$, where $\bar{d}_\alpha := \lim_{i \rightarrow +\infty} \mathbb{E}[d_\alpha(i)]$. In order to maintain stability of the system we must require

$$\frac{1}{n} \sum_{\alpha=1}^m \lambda_\alpha \left(\frac{r_\alpha \bar{d}_\alpha}{v} + \bar{s}_\alpha \right) < 1. \quad (6)$$

Applying the definition of ϱ in Eq. (2), we write Eq. (6) as

$$\sum_{\alpha=1}^m r_\alpha \lambda_\alpha \bar{d}_\alpha < (1 - \varrho)nv. \quad (7)$$

For a stable policy P , let \bar{N}_α represent the steady-state expected number of unserved α -demands. Then, the expected total number of outstanding demands that require on-site service (i.e., cannot be serviced remotely) is given by $\sum_{j=1}^m r_j \bar{N}_j$. We now apply a result from the dynamic traveling repairperson problem (see [12], page 23) which states that in heavy load ($\varrho \rightarrow 1^-$), if the steady-state number of outstanding demands is N , then a lower bound on expected travel distance between demands is $(\beta_{\text{TSP}}/\sqrt{2})\sqrt{|\mathcal{E}|/N}$. Applying this result we have that

$$\bar{d}_\alpha \geq \frac{\beta_{\text{TSP}}}{\sqrt{2}} \sqrt{\frac{|\mathcal{E}|}{\sum_j r_j \bar{N}_j}} =: \bar{d}, \quad (8)$$

for each $\alpha \in \{1, \dots, m\}$. Combining with Eq. (7), squaring both sides, and rearranging we obtain

$$\frac{\beta_{\text{TSP}}^2}{2} \frac{|\mathcal{E}| (\sum_\alpha r_\alpha \lambda_\alpha)^2}{n^2 v^2 (1 - \varrho)^2} < \sum_\alpha r_\alpha \bar{N}_\alpha.$$

From Little's law, $\bar{N}_\alpha = \lambda_\alpha W_\alpha$ for each $\alpha \in \{1, \dots, m\}$, and thus

$$\sum_\alpha r_\alpha \lambda_\alpha W_\alpha > \frac{\beta_{\text{TSP}}^2}{2} \frac{|\mathcal{E}|}{n^2 v^2 (1 - \varrho)^2} \left(\sum_\alpha r_\alpha \lambda_\alpha \right)^2. \quad (9)$$

Recalling that $W_\alpha = D_\alpha - \bar{s}_\alpha$ and $r_\alpha \in \{0, 1\}$ for each $\alpha \in \{1, \dots, m\}$, we see that Eq. (9) gives us $2^m - 1$ constraints

on the feasible values of $D_1(P), \dots, D_m(P)$. Hence, a lower bound on D^* can be found by minimizing $\sum_{\alpha=1}^m W_\alpha$ subject to the constraints in Eq. (9). By considering the dual of this problem, one can verify that under the class labeling in Eq. (3), the problem is equivalent to:

$$\text{minimize} \quad \sum_{\alpha=1}^m c_\alpha W_\alpha,$$

subject to

$$\begin{bmatrix} \lambda_1 & 0 & 0 & \cdots & 0 \\ \lambda_1 & \lambda_2 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & & 0 \\ \lambda_1 & \lambda_2 & \lambda_3 & \cdots & \lambda_m \end{bmatrix} \begin{bmatrix} W_1 \\ W_2 \\ \vdots \\ W_m \end{bmatrix} \geq \Psi \begin{bmatrix} \lambda_1^2 \\ (\lambda_1 + \lambda_2)^2 \\ \vdots \\ (\lambda_1 + \cdots + \lambda_m)^2 \end{bmatrix},$$

where

$$\Psi := \frac{\beta_{\text{TSP}}^2}{2} \frac{|\mathcal{E}|}{n^2 v^2 (1 - \varrho)^2}.$$

Under the class labeling in Eq. (3) the above linear program is feasible and bounded, and its solution (W_1^*, \dots, W_m^*) is given by

$$W_\alpha^* = \Psi \left(\lambda_\alpha + 2 \sum_{j=1}^{\alpha-1} \lambda_j \right).$$

After rearranging, the optimal value of the cost function, and thus the lower bound on D^* , is given by

$$\sum_{\alpha=1}^m c_\alpha W_\alpha^* = \Psi \sum_{\alpha=1}^m \left(c_\alpha + 2 \sum_{j=\alpha+1}^m c_j \right) \lambda_\alpha.$$

Applying the definition of Ψ we obtain the desired result. ■

IV. SEPARATE QUEUES POLICY

In this section we introduce and analyze the Separate Queues (SQ) policy. We show that this policy is within a factor of $2m^2$ of the lower bound in heavy load.

To present the SQ policy we need some notation. We assume vehicle $k \in \{1, \dots, n\}$ has a service region $R^{[k]} \subset \mathcal{E}$, such that $\{R^{[1]}, \dots, R^{[n]}\}$ form a partition of the environment \mathcal{E} . In general, the partition could be time varying, but for the description of the SQ policy this will not be required. We assume that information on outstanding demands of type $\alpha \in \{1, \dots, m\}$ in region $R^{[k]}$ at time t is summarized as a finite set of demand positions $Q_\alpha^{[k]}(t)$ with $N_\alpha^{[k]}(t) := \text{card}(Q_\alpha^{[k]}(t))$. Demands of type α with location in $R^{[k]}$ are inserted in the set $Q_\alpha^{[k]}$ as soon as they are generated. Removal from the set $Q_\alpha^{[k]}$ requires that service vehicle k moves to the demand location, and provides the on-site service. With this notation the policy is given as Algorithm 1.

A. Stability Analysis of the SQ Policy in Heavy Load

In this section we analyze the SQ policy in heavy load, i.e., as $\varrho \rightarrow 1^-$. In the SQ policy each region $R^{[k]}$ has equal area, and contains a single vehicle. Thus, the n vehicle problem in a region of area $|\mathcal{E}|$ has been turned into n independent single-vehicle problems, each in a region of area $|\mathcal{E}|/n$, with arrival rates λ_α/n . To determine the performance of the

Algorithm 1: Separate Queues (SQ) Policy

Assumes: A probability distribution $\mathbf{p} = [p_1, \dots, p_m]$.

- 1 Partition \mathcal{E} into n equal area regions and assign one vehicle to each region.
 - 2 **foreach** vehicle-region pair k **do**
 - 3 **if** the set $\cup_{\alpha} Q_{\alpha}^{[k]}$ is empty **then**
 - 4 Move vehicle toward the median of its own region until a demand arrives.
 - 5 **else**
 - 6 Select $Q \in \{Q_1^{[k]}, \dots, Q_m^{[k]}\}$ according to \mathbf{p} .
 - 7 **if** Q is empty **then**
 - 8 Reselect until Q is nonempty.
 - 9 Compute TSP tour through all demands in Q .
 - 10 Service Q following the TSP tour, starting at the demand closest to the vehicle's current position.
 - 11 Repeat.
 - 12 Optimize over \mathbf{p} .
-

policy we need only study the performance in a single region k . For simplicity of notation we omit the label k . We refer to the time instant t_i in which the vehicle computes a new TSP tour as the epoch i of the policy; we refer to the time interval between epoch i and epoch $i+1$ as the i th iteration and we will refer to its length as T_i . Finally, let $N_{\alpha}(t_i) := N_{\alpha,i}$, $\alpha \in \{1, \dots, m\}$, be the number of outstanding α -demands at beginning of iteration i .

The following straightforward lemma, proved in [9], will be essential in deriving our main results.

Lemma IV.1 (Number of outstanding demands) In heavy load (i.e., $\varrho \rightarrow 1^-$), after a transient, the number of demands serviced in a single tour of the vehicle in the SQ policy is very large with high probability (i.e., the number of demands tends to $+\infty$ with probability that tends to 1, as ϱ approaches 1^-).

Let TS_j be the event that Q_j is selected for service at iteration i of the SQ policy. By the law of total probability

$$\mathbb{E}[N_{\alpha,i+1}] = \sum_{j=1}^m p_j \mathbb{E}(N_{\alpha,i+1}|TS_j), \quad \alpha \in \{1, \dots, m\},$$

where the conditioning is with respect to the task being performed during iteration i . During iteration i of the policy, demands arrive according to independent Poisson processes. Call $N_{\alpha,i}^{\text{new}}$ the α -demands ($\alpha \in \{1, \dots, m\}$) newly arrived during iteration i ; then, by definition of the SQ policy

$$\begin{aligned} \mathbb{E}(N_{\alpha,i+1}|TS_j) &= \begin{cases} \mathbb{E}(N_{\alpha,i}^{\text{new}}|TS_j), & \text{if } \alpha = j \\ \mathbb{E}(N_{\alpha,i}|TS_j) + \mathbb{E}(N_{\alpha,i}^{\text{new}}|TS_j), & \text{o.w.} \end{cases} \end{aligned}$$

By the law of iterated expectation, we have $\mathbb{E}(N_{\alpha,i}^{\text{new}}|TS_j) = (\lambda_{\alpha}/n)\mathbb{E}(T_i|TS_j)$. Moreover, since the number of demands outstanding at the beginning of

iteration i is independent of the task that will be chosen, we have $\mathbb{E}(N_{\alpha,i}|TS_j) = \mathbb{E}[N_{\alpha,i}]$. Thus we obtain

$$\mathbb{E}(N_{\alpha,i+1}|TS_j) = \begin{cases} \frac{\lambda_{\alpha}}{n}\mathbb{E}(T_i|TS_j), & \text{if } \alpha = j \\ \mathbb{E}[N_{\alpha,i}] + \frac{\lambda_{\alpha}}{n}\mathbb{E}(T_i|TS_j), & \text{o.w.} \end{cases}$$

Therefore, we are left with computing the conditional expected values of T_i . The length of T_i is given by the time needed by the vehicle to travel along the TSP tour plus the time spent to service demands. Assuming i large enough, Lemma (IV.1) holds, and we can apply Eq. (1) to estimate from the quantities $N_{\alpha,i}$, $\alpha \in \{1, \dots, m\}$, the length of the TSP tour at iteration i . Conditioning on TS_j (when only demands of type j are serviced), we have

$$\begin{aligned} \mathbb{E}(T_i|TS_j) &= \frac{\beta_{\text{TSP}}\sqrt{|\mathcal{E}|/n}}{v} \mathbb{E}(\sqrt{N_{j,i}}|TS_j) \\ &\quad + \mathbb{E}\left(\sum_{k=1}^{N_{j,i}} s_{j,k}|TS_j\right) \\ &\leq \frac{\beta_{\text{TSP}}\sqrt{|\mathcal{E}|/n}}{v} \sqrt{\mathbb{E}[N_{j,i}]} + \mathbb{E}[N_{j,i}]\bar{s}_j, \end{aligned}$$

where we have: (i) applied Eq. (1), (ii) applied Jensen's inequality for concave functions, in the form $\mathbb{E}[\sqrt{X}] \leq \sqrt{\mathbb{E}[X]}$, (iii) removed the conditioning on TS_j , since the random variables $N_{\alpha,i}$ are independent from *future* events, and in particular from the choice of the task at iteration i , and (iv) used the fact that the on-site service times are independent from the number of outstanding demands.

Collecting the above results (and using the shorthand \bar{X} to indicate $\mathbb{E}[X]$, where X is any random variable), we have

$$\begin{aligned} \bar{N}_{\alpha,i+1} &\leq (1-p_{\alpha})\bar{N}_{\alpha,i} \\ &\quad + \sum_{j=1}^m p_j \frac{\lambda_{\alpha}}{n} \left[\frac{\beta_{\text{TSP}}\sqrt{|\mathcal{E}|/n}}{v} \sqrt{\bar{N}_{j,i}} + \bar{N}_{j,i}\bar{s}_j \right], \end{aligned} \quad (10)$$

for each $\alpha \in \{1, \dots, m\}$. The m inequalities above give a system of recursive relations that describe an upper bound on $\bar{N}_{\alpha,i}$, $\alpha \in \{1, \dots, m\}$. The following theorem bounds the values to which they converge.

Theorem IV.2 (Queue length) In heavy load, for every set of initial conditions $\{\bar{N}_{\alpha,0}\}_{\alpha \in \{1, \dots, m\}}$, the trajectories $i \rightarrow \bar{N}_{\alpha,i}$, $\alpha \in \{1, \dots, m\}$, resulting from Eqs. (10), satisfy

$$\limsup_{i \rightarrow +\infty} \bar{N}_{\alpha,i} \leq \frac{\beta_{\text{TSP}}^2 |\mathcal{E}|}{n^3 v^2 (1-\varrho)^2} \frac{\lambda_{\alpha}}{p_{\alpha}} \left(\sum_{j=1}^m \sqrt{\lambda_j p_j} \right)^2.$$

Due to space constraints, the proof is omitted and can be found in [13].

B. Delay of the SQ Policy in Heavy Load

From Theorem IV.2, and using Little's law, the delay of α -demands is

$$\begin{aligned} D_{\alpha}(\text{SQ}) &\leq \frac{n}{\lambda_{\alpha}} \limsup_{i \rightarrow +\infty} \bar{N}_{\alpha,i} + \bar{s}_{\alpha} \\ &= \frac{\beta_{\text{TSP}}^2 |\mathcal{E}|}{n^2 v^2 (1-\varrho)^2} \frac{1}{p_{\alpha}} \left(\sum_{j=1}^m \sqrt{\lambda_j p_j} \right)^2, \end{aligned}$$

where we neglected \bar{s}_α because of the heavy-load assumption.

Thus, the delay (as defined in Eq. (4)) of the SQ policy, satisfies in heavy load

$$D(\text{SQ}) \leq \frac{\beta_{\text{TSP}}^2 |\mathcal{E}|}{n^2 v^2 (1 - \varrho)^2} \sum_{\alpha=1}^m \frac{c_\alpha}{p_\alpha} \left(\sum_{i=1}^m \sqrt{\lambda_i p_i} \right)^2. \quad (11)$$

With this expression we prove our main result on the performance of the SQ policy.

Theorem IV.3 (SQ policy performance) In heavy load, the delay of the SQ policy is within a factor $2m^2$ of the optimal, independent of the arrival rates $\lambda_1, \dots, \lambda_m$, coefficients c_1, \dots, c_m , service times $\bar{s}_1, \dots, \bar{s}_m$, and the number of vehicles n .

Proof: We would like to compare the performance of this policy with the lower bound. To do this, consider setting $p_\alpha := c_\alpha$ for each $\alpha \in \{1, \dots, m\}$. Defining $B := \beta_{\text{TSP}}^2 |\mathcal{E}| / (n^2 v^2 (1 - \varrho)^2)$, Eq. (11) can be written as

$$D(\text{SQ}) \leq Bm \left(\sum_{i=1}^m \sqrt{c_i \lambda_i} \right)^2.$$

Next, the lower bound in Eq. (5) is

$$D^* \geq \frac{B}{2} \sum_{i=1}^m \left(c_i + 2 \sum_{j=i+1}^m c_j \right) \lambda_i \geq \frac{B}{2} \sum_{i=1}^m (c_i \lambda_i).$$

Thus, comparing the upper and lower bounds

$$\frac{D(\text{SQ})}{D^*} \leq 2m \frac{\left(\sum_{i=1}^m \sqrt{c_i \lambda_i} \right)^2}{\sum_{i=1}^m (c_i \lambda_i)}. \quad (12)$$

Letting $x_i := \sqrt{c_i \lambda_i}$, and $\mathbf{x} := [x_1, \dots, x_m]$, the numerator of the fraction in Eq. (12) is $\|\mathbf{x}\|_1^2$, and the denominator is $\|\mathbf{x}\|_2^2$. But the one- and two-norms of a vector $\mathbf{x} \in \mathbb{R}^m$ satisfy $\|\mathbf{x}\|_1 \leq \sqrt{m} \|\mathbf{x}\|_2$. Thus, in heavy load we obtain

$$\frac{D(\text{SQ})}{D^*} \leq 2m \left(\frac{\|\mathbf{x}\|_1}{\|\mathbf{x}\|_2} \right)^2 \leq 2m^2,$$

and the policy is a $2m^2$ -factor approximation. ■

Remark IV.4 (Relation to RP policy in [9]) For $m = 2$ the SQ policy is within a factor of 8 of the optimal. This improves on the factor of 12 obtained for the Randomized Priority (RP) policy in [9]. However, it appears that the RP policy bound is not tight, since for two classes, simulations indicate it performs no worse than the SQ policy. □

V. SIMULATIONS AND DISCUSSION

In this section we discuss, through the use of simulations, the performance of the SQ policy with the probability assignment $p_\alpha := c_\alpha$, $\alpha \in \{1, \dots, m\}$. In particular, we study (i) conditions for which the gap between lower bound in Eq. (5) and upper bound in Eq. (11) is maximized, (ii) the suboptimality of the probability assignment $p_\alpha = c_\alpha$, and, finally, (iii) how different the cost function in Eq. (4) may

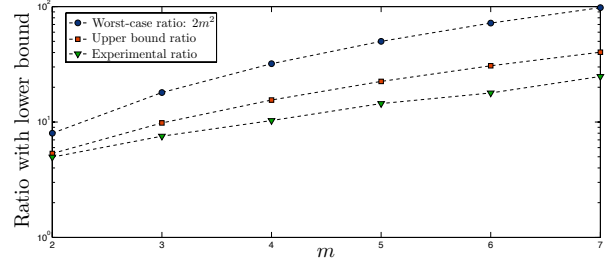


Fig. 1. Experimental results for the SQ policy in worst-case conditions; $\varrho = 0.85$ and $\lambda_1 = 1$.

be, in general, for the SQ policy and a policy that services demands all together irrespective of priorities. Simulations of the SQ policy were performed using `linkern`¹ as a solver to generate approximations to the optimal TSP tour.

A. Unfavorable Conditions for the SQ Policy

One may question if for some sets $\{\lambda_\alpha\}$ and $\{c_\alpha\}$, $\alpha \in \{1, \dots, m\}$, the ratio between upper bound (11) and lower bound (5) is indeed close to $2m^2$. The answer is affirmative: consider, e.g., the case $\lambda_1 \ll \lambda_2 \ll \dots \ll \lambda_m$ and $c_1 \gg c_2 \gg \dots \gg c_m$, with $\lambda_\alpha c_\alpha = a$, for some positive constant a . Then, the upper bound is equal to $Bm^3 a$ and the lower bound is approximately equal to $Bma/2$, thus their ratio is (arbitrarily) close to $2m^2$. Then, we simulated the SQ policy for the case $\lambda_m = b\lambda_{m-1} = b^2\lambda_{m-2} = \dots = b^{m-1}\lambda_1$ and $c_1 = bc_2 = \dots = b^{m-1}c_m$ with $b = 2$. Fig. 1 shows that the experimental value of the cost function (averaged over 10 simulation runs) indeed increases proportionally to m^2 .

B. Suboptimality of the Approximate Probability Assignment

To prove Theorem IV.3 we used the probability assignment

$$p_\alpha := c_\alpha \quad \text{for each } \alpha \in \{1, \dots, m\}. \quad (13)$$

However, one would like to select $[p_1, \dots, p_m] =: \mathbf{p}$ that minimizes the right-hand side of Eq. (11). The minimization of the right-hand side of Eq. (11) is a constrained multi-variable nonlinear optimization problem over \mathbf{p} , that is, in m dimensions. However, for two classes of demands the optimization is over a single variable p_1 , and it can be readily solved. A comparison of optimized upper bound, denoted upbd_{opt} , with the upper bound obtained using the probability assignment in Eq. (13), denoted upbd_c , is shown in Fig. 2.

For $m > 2$ we approximate the solution of the optimization problem as follows. For each value of m we perform 1000 runs. In each run we randomly generate $\lambda_1, \dots, \lambda_m$, c_1, \dots, c_m , and five sets of initial probability assignments $\mathbf{p}_1, \dots, \mathbf{p}_5$. From each initial probability assignment we use a line search to locally optimize the probability assignment. We take the ratio between upbd_c and the least upper bound $\text{upbd}_{\text{local opt}}$ obtained from the five locally optimized probability assignments. We also record the maximum variation in

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

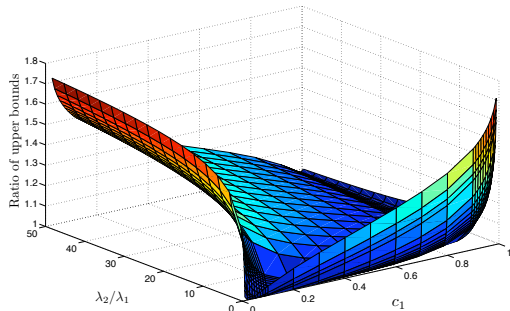


Fig. 2. The ratios $\text{upbd}_c/\text{upbd}_{\text{opt}}$ for 2 classes of demands.

Number of classes (m)	$\max \text{upbd}_c/\text{upbd}_{\text{opt}}$	Max. % variation
3	1.60	0.12
4	1.51	0.04
5	1.51	0.08
6	1.74	0.02
7	1.88	0.08
8	1.63	0.15

TABLE I
RATIO OF UPPER BOUND WITH $p_\alpha = c_\alpha$ AND UPPER BOUND WITH OPTIMIZED \mathbf{p} .

the five locally optimized upper bounds. This is summarized in Table I. The second column shows the largest ratio obtained over the 1000 runs. The third column shows the largest % variation in the 1000 runs. The assignment in Eq. (13) performs within a factor of two of the optimized assignment. In addition, the optimization appears to converge to values close to a global optimum since all five random conditions converge to values that are within $\sim 0.1\%$ of each other on every run.

C. The Merge Policy

The simplest possible policy for our problem would be to ignore priorities and service demands all together, by repeatedly forming TSP tours of outstanding demands (i.e., by using the SQ policy as though there were only one class). We call such a policy the Merge policy. However, the performance of the SQ and the Merge policy can be arbitrarily far apart. Indeed, by defining the overall arrival rate $\Lambda := \sum_{\alpha=1}^m \lambda_\alpha$ and overall mean on-site service $\bar{S} := \sum_{\alpha=1}^m \bar{s}_\alpha$, and by using the upper bounds in [4], we immediately obtain as an upper bound for the Merge policy: $D(\text{Merge}) \leq \frac{\beta_{\text{TSP}}^2 |\mathcal{E}| \Lambda}{n^2 v^2 (1-\rho)^2}$. Then, we see that $D(\text{Merge})/D(\text{SQ})$ can be arbitrarily large by choosing $\lambda_m \gg \lambda_\alpha$ and $c_m \ll c_\alpha$, with $\alpha \in \{1, \dots, m-1\}$. This behavior is confirmed by experimental results, as depicted in Fig. 3 where we show the experimental ratios of delays between Merge and SQ policy (the ratios are averaged values over 10 simulation runs).

VI. CONCLUSIONS

In this paper we studied a dynamic multi-vehicle routing problem with multiple classes of demands. For every set of coefficients, we determined a lower bound on the achievable convex combination of the class delays. We presented the

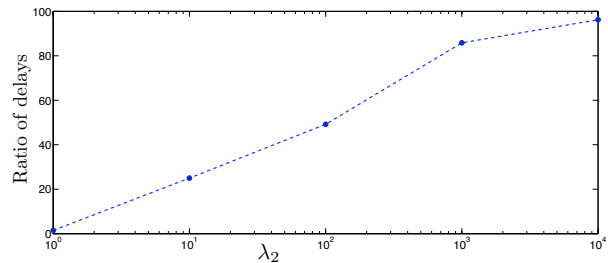


Fig. 3. Ratio of experimental delays between Merge policy and SQ policy as a function of λ_2 , with $m = 2$, $\lambda_1 = 1$, $c = 0.995$ and $\rho = 0.9$.

Separate Queues (SQ) policy and showed that its deviation from the lower bound depends only on the number of the classes. We believe that there is room for improvement in the lower bound, and thus the SQ policy's performance may be significantly better than is indicated by its deviation from the current lower bound. Thus, our main thrust of future work will be in trying to raise the lower bound. We are also interested in combining the aspects of multi-class vehicle routing with problems in which demands require teams of vehicles for their service, and in extending our results to the case of non-uniform demand densities (possibly class dependent).

REFERENCES

- [1] L. Kleinrock, *Queueing Systems. Volume II: Computer Applications*. New York, NY: John Wiley and Sons, 1976.
- [2] E. G. Coffman Jr. and I. Mitrani, "A characterization of waiting time performance realizable by single-server queues," *Operations Research*, vol. 28, no. 3, pp. 810–821, 1980.
- [3] D. Bertsimas, I. C. Paschalidis, and J. N. Tsitsiklis, "Optimization of multiclass queueing networks: Polyhedral and nonlinear characterizations of achievable performance," *The Annals of Applied Probability*, vol. 4, no. 1, pp. 43–75, 1994.
- [4] 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.
- [5] —, "Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles," *Operations Research*, vol. 41, no. 1, pp. 60–76, 1993.
- [6] E. Frazzoli and F. Bullo, "Decentralized algorithms for vehicle routing in a stochastic time-varying environment," in *Proc CDC*, Paradise Island, Bahamas, Dec. 2004, pp. 3357–3363.
- [7] M. Pavone, E. Frazzoli, and F. Bullo, "Decentralized algorithms for stochastic and dynamic vehicle routing with general target distribution," in *Proc CDC*, New Orleans, LA, Dec. 2007, pp. 4869–4874.
- [8] R. C. Larson and A. R. Odoni, *Urban Operations Research*. Prentice Hall, 1981.
- [9] S. L. Smith, M. Pavone, F. Bullo, and E. Frazzoli, "Dynamic vehicle routing with heterogeneous demands," in *Proc CDC*, Cancún, México, Dec. 2008, pp. 1206–1211.
- [10] J. M. Steele, "Probabilistic and worst case analyses of classical problems of combinatorial optimization in Euclidean space," *Mathematics of Operations Research*, vol. 15, no. 4, p. 749, 1990.
- [11] 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.
- [12] H. Xu, "Optimal policies for stochastic and dynamic vehicle routing problems," Dept. of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA, 1995.
- [13] M. Pavone, S. L. Smith, F. Bullo, and E. Frazzoli, "Dynamic multi-vehicle routing with multiple classes of demands," Mar. 2009, available at <http://arXiv.org/abs/0903.2695>.