

MIT Open Access Articles

*Dynamic Vehicle Routing with Priority
Classes of Stochastic Demands*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Smith, Stephen L. et al. "Dynamic Vehicle Routing with Priority Classes of Stochastic Demands." SIAM Journal on Control and Optimization 48.5 (2010): 3224-3245. ©2010 Society for Industrial and Applied Mathematics

As Published: <http://dx.doi.org/10.1137/090749347>

Publisher: Society of Industrial and Applied Mathematics

Persistent URL: <http://hdl.handle.net/1721.1/57470>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



DYNAMIC VEHICLE ROUTING WITH PRIORITY CLASSES OF STOCHASTIC DEMANDS*

STEPHEN L. SMITH[†], MARCO PAVONE[‡], FRANCESCO BULLO[§], AND
EMILIO FRAZZOLI[‡]

Abstract. In this paper we introduce a dynamic vehicle routing problem in which there are multiple vehicles and multiple priority classes of service demands. Service demands of each priority 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, 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.

Key words. vehicle routing, optimization, multiagent systems, heterogeneous systems

AMS subject classifications. 68T40, 90C27, 90B22, 93E03, 37N35

DOI. 10.1137/090749347

1. Introduction. A classic problem in queueing theory is that of priority queueing [10]. In the simplest setup, customers arrive at a single server sequentially over time. Each customer is a member of either the high-priority or the low-priority class. High-priority customers and low-priority customers form separate queues. The goal is to provide the highest possible quality of service to the high-priority queue (Q_1) while maintaining stability of the low-priority queue (Q_2). That is, the goal is to minimize the expected delay for high-priority customers while keeping the length of the low-priority queue finite. When both the customer interarrival times and the customer service times are distributed exponentially, the preemptive priority policy is known to be optimal [10]:

When Q_1 is nonempty, serve high-priority customers; when Q_1 is empty, serve low-priority customers. If a high-priority customer arrives while serving Q_2 , then preempt service and immediately begin serving the high-priority customer.

*Received by the editors February 11, 2009; accepted for publication (in revised form) October 19, 2009; published electronically January 15, 2010. 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. Preliminary versions of this work appeared in [21] and [17].

<http://www.siam.org/journals/sicon/48-5/74934.html>

[†]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 (slsmith@mit.edu).

[‡]Aeronautics and Astronautics Department, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139 (pavone@mit.edu, frazzoli@mit.edu).

[§]Department of Mechanical Engineering, Center for Control, Dynamical Systems and Computation, University of California, Santa Barbara, CA 93106 (bullo@engineering.ucsb.edu).

A more general two-class queueing problem is to minimize a convex combination of the service delays for high- and low-priority customers:

$$cD_1 + (1 - c)D_2, \quad \text{where } c \in (0, 1).$$

In this case an optimal policy can be created by using a mixed policy that spends fraction c of the time serving Q_1 as the high-priority queue and fraction $(1 - c)$ serving Q_2 as though it is the high-priority queue [8]. The set of achievable delays has also been studied in the more general setting of queueing networks [2].

In this paper we consider an m -class, n -service-vehicle spatial queueing problem, called *dynamic vehicle routing with priority classes*. Demands for service arrive sequentially over time in a compact environment \mathcal{E} in the plane. Each demand is a member of one of m priority classes. Upon arrival, each 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 demand's arrival and its service completion. Then, given convex combination coefficients $c_1, \dots, c_m > 0$, the goal is to find the vehicle routing policy that minimizes $c_1D_1 + \dots + c_mD_m$. By increasing the coefficients for certain classes, a higher-priority level can be given to their demands. This problem has important applications in areas such as UAV surveillance, where targets are given different priority levels based on their urgency or potential importance [1].

When there is only one class of demands, the problem in this paper is known as the dynamic traveling repairperson problem (DTRP) [19, 3, 4]. The DTRP was arguably the first member of a class of problems known as dynamic vehicle routing (DVR); for other DVR problems see [11]. In [3, 4], optimal DTRP policies are proposed in 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 [14], a simple unified DTRP policy is presented for both light- and heavy-load conditions. Recently, there has been an increased interest in DVR among researchers in robotic motion planning, as it provides a powerful method for completing spatially distributed tasks that are generated in real-time. In particular, DVR results have recently been obtained on decentralized policies [9, 16], moving demands [7], impatient demands [15], demands requiring pickup and delivery [24], and demands which require teams of vehicles [20]. Other related vehicle routing problems include the orienteering problem [6], the minimum latency problem [5], the dynamic assignment problem [22], and spatial queueing in the context of urban operations research [12].

The main contribution of this paper is to introduce DVR with priority classes. We derive a lower bound on the achievable values of the convex combination of delays and propose a novel policy in which each class of demands is served separately from the others. We show that in heavy-load the policy performs within a constant factor $2m^2$ of the lower bound. 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. To establish the constant factor, we proceed in a manner similar to [14, 15] and develop a system of nonlinear inequality-based recursive equations on the expected number of outstanding demands. We then utilize a novel proof technique to upper bound the limiting number of outstanding demands, which relies on constructing a set of linear equality-based recursive equations to bound trajectories. We present an

improvement on the policy in which classes of similar priority are merged together. We also perform extensive simulations and introduce an effective heuristic improvement called the tube heuristic.

The paper is organized as follows. In section 2 we give some asymptotic properties of optimal Euclidean traveling salesperson tours. In section 2.2 we formalize the problem, and in section 3 we derive a lower bound on the achievable delay. In section 4 we introduce and analyze the separate queues policy and present the improvements given by *queue merging* and the *tube heuristic*. Finally, in section 5 we present simulation results.

2. Background and problem statement. In this section we summarize the asymptotic properties of optimal Euclidean traveling salesperson tours and formalize the DVR problem with priority classes.

2.1. 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 [23] it is shown that there exists a constant β_{TSP} such that, almost surely,

$$(2.1) \quad \lim_{N \rightarrow +\infty} \frac{\text{TSP}(Q)}{\sqrt{N}} = \beta_{\text{TSP}} \sqrt{|\mathcal{E}|}.$$

The constant β_{TSP} has been estimated numerically as $\beta_{\text{TSP}} \approx 0.7120 \pm 0.0002$ [18]. The limit in (2.1) holds for all compact sets \mathcal{E} , and the shape of \mathcal{E} affects only the convergence rate to the limit. In [12], the authors note that if \mathcal{E} is “fairly compact [square] and fairly convex”, then (2.1) provides an adequate estimate of the optimal TSP tour length for values of N as low as 15.

2.2. 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} ; see Figure 2.1. 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

$$(2.2) \quad \varrho := \frac{1}{n} \sum_{\alpha=1}^m \lambda_\alpha \bar{s}_\alpha < 1.$$

The *load factor* ϱ is a standard quantity in queueing theory [10] 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 $\varrho \in [0, 1)$, and a




FIG. 2.1. A depiction of the problem for two vehicles and three priority classes. Left figure: One vehicle is moving to a class 1 demand and the other to a class 2 demand. Right figure: The bottom vehicle has serviced the class 1 demand and is moving to a class 2 demand. A new class 3 demand has arrived.

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 is defined as $D_\alpha(P) := \lim_{i \rightarrow +\infty} \mathbb{E}[D_\alpha(i)]$, and the steady-state wait time 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

$$(2.3) \quad \frac{c_1}{\lambda_1} \geq \frac{c_2}{\lambda_2} \geq \dots \geq \frac{c_m}{\lambda_m},$$

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 DVR with priority classes. Given the coefficients $c_\alpha > 0$, $\alpha \in \{1, \dots, m\}$, with $\sum_{\alpha=1}^m c_\alpha = 1$, and satisfying (2.3), let $D(P) := \sum_{\alpha=1}^m c_\alpha D_\alpha(P)$ be the cost of policy $P \in \Pi$. Then, the problem is to determine a vehicle routing policy P^* , if one exists, such that

$$(2.4) \quad D(P^*) = \inf_{P \in \Pi} D(P).$$

We let D^* denote the right-hand side of (2.4). 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. In the light-load regime, existing policies for the dynamic traveling repairperson can be used, as is summarized in the following remark.

Remark 2.1 (light-load regime). In light-load, $\rho \rightarrow 0^+$, optimal policies have been developed for the DTRP (i.e., the single-class DVR problem). These policies rely on the computation of a set of n -median locations for the environment \mathcal{E} , that is, a set of n positions $\mathcal{Q}^* \subset \mathcal{E}$, that minimize

$$\mathbb{E} \left[\min_{\mathbf{q} \in \mathcal{Q}^*} \|\mathbf{q} - \mathbf{q}_0\| \right],$$

where \mathbf{q}_0 is a uniformly distributed location in \mathcal{E} . In particular, the n stochastic queue median (n SQM) policy, first introduced in [4], can be described as follows:

Place one vehicle at each of the n -median locations of \mathcal{E} . When a demand arrives, assign it to the closest median location and to the corresponding vehicle. Have each vehicle service its demands in the first-come-first-served order, returning to its median location after each service is completed.

In fact, by following the proof in [4], one can show that the n SQM policy is an optimal policy for DVR with priority classes. The proof of this statement is omitted in the interest of brevity, and we refer interested readers to [4] for details.

3. Lower bound in heavy-load. In this section we present two lower bounds on the delay in (2.4). The first holds only in heavy-load (i.e., as $\rho \rightarrow 1^-$), while the second (less tight) bound holds for all ρ .

THEOREM 3.1 (heavy-load lower bound). *For every routing policy P ,*

$$(3.1) \quad D(P) \geq \frac{\beta_{\text{TSP}}^2 |\mathcal{E}|}{2n^2 v^2 (1-\rho)^2} \sum_{\alpha=1}^m \left(c_\alpha + 2 \sum_{j=\alpha+1}^m c_j \right) \lambda_\alpha \quad \text{as } \rho \rightarrow 1^-,$$

where c_1, \dots, c_m satisfy (2.3).

Before proving Theorem 3.1 let us quickly comment on the form of (3.1). The right-hand side of (3.1) approaches $+\infty$ as $\rho \rightarrow 1^-$. Thus, one should more formally write the inequality with $D(P)(1-\rho)^2$ on the left-hand side, so that the right-hand side is finite. However, this makes the presentation less readable, and thus henceforth we adhere to the less formal but more transparent style of (3.1).

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 $r_\alpha d_\alpha(i)/v + s_\alpha(i)$, where v is the service vehicle speed. The steady-state expected service requirement

is $r_\alpha \bar{d}_\alpha / v + s_\alpha$, where $\bar{d}_\alpha := \lim_{i \rightarrow +\infty} \mathbb{E}[d_\alpha(i)]$. In order to maintain stability of the system we must require

$$(3.2) \quad \frac{1}{n} \sum_{\alpha=1}^m \lambda_\alpha \left(\frac{r_\alpha \bar{d}_\alpha}{v} + \bar{s}_\alpha \right) < 1.$$

Applying the definition of ρ in (2.2), we write (3.2) as

$$(3.3) \quad \sum_{\alpha=1}^m r_\alpha \lambda_\alpha \bar{d}_\alpha < (1 - \rho)nv.$$

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 DTRP (see [26, page 23]), which states that, in heavy-load ($\rho \rightarrow 1^-$), 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

$$(3.4) \quad \bar{d}_\alpha \geq \frac{\beta_{\text{TSP}}}{\sqrt{2}} \sqrt{\frac{|\mathcal{E}|}{\sum_j r_j \bar{N}_j}} =: \bar{d}$$

for each $\alpha \in \{1, \dots, m\}$. Combining (3.3) and (3.4),

$$\frac{\sum_\alpha r_\alpha \lambda_\alpha}{nv(1 - \rho)} < \frac{1}{\bar{d}}.$$

Applying the definition of \bar{d} , squaring both sides, and rearranging we obtain

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

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

$$(3.5) \quad \sum_\alpha r_\alpha \lambda_\alpha W_\alpha > \frac{\beta_{\text{TSP}}^2}{2} \frac{|\mathcal{E}|}{n^2 v^2 (1 - \rho)^2} \left(\sum_\alpha r_\alpha \lambda_\alpha \right)^2.$$

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 (3.5) 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 $\sum_{\alpha=1}^m c_\alpha W_\alpha$ subject to the constraints in (3.5). We can lower bound the solution to the optimization problem by minimizing the cost function subject to only a subset of the $2^m - 1$ constraints. In particular, consider the following linear program:

$$\begin{aligned} & \text{minimize} && \sum_{\alpha=1}^m c_\alpha W_\alpha, \\ & \text{subject to} && \begin{bmatrix} \lambda_1 & 0 & 0 & \cdots & 0 \\ \lambda_1 & \lambda_2 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & & 0 \\ \lambda_1 & \lambda_2 & \lambda_3 & \cdots & \lambda_m \end{bmatrix} \begin{bmatrix} W_1 \\ W_2 \\ \vdots \\ W_m \end{bmatrix} \geq \Psi \begin{bmatrix} \lambda_1^2 \\ (\lambda_1 + \lambda_2)^2 \\ \vdots \\ (\lambda_1 + \cdots + \lambda_m)^2 \end{bmatrix}, \end{aligned}$$




FIG. 3.1. The feasible region of the linear program for two queues. When class 1 is of higher priority, the solution is given by the corner. Otherwise, the solution is $-\infty$.

where

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

The above problem is feasible (see Figure 3.1), it has only one basic feasible solution, and it is of the following form: minimize $\mathbf{c}^T \mathbf{W}$ subject to $\mathbf{A}\mathbf{W} \geq \mathbf{b}$. Thus, either the problem is unbounded, or the solution \mathbf{W}^* is given by the basic feasible solution. To establish boundedness we consider the following dual problem: maximize $\mathbf{b}^T \mathbf{y}$ subject to $A^T \mathbf{y} = \mathbf{c}$ and $\mathbf{y} \geq 0$. By the duality theorem of linear programming [13], if the dual is feasible, then the minimization problem is bounded. To check feasibility of the dual, we solve for $A^T \mathbf{y} = \mathbf{c}$, with $\mathbf{y} \geq 0$, to obtain

$$y_\alpha = \frac{c_\alpha}{\lambda_\alpha} - \frac{c_{\alpha+1}}{\lambda_{\alpha+1}} \geq 0 \quad \text{for all } \alpha \in \{1, \dots, m-1\},$$

$$y_m = \frac{c_m}{\lambda_m} \geq 0.$$

Thus, the dual is feasible if and only if the priority classes are labeled as in (2.3). When (2.3) is satisfied, the minimization problem is bounded, and its solution (W_1^*, \dots, W_m^*) is given by

$$W_\alpha^* = \frac{\Psi}{\lambda_\alpha} \left((\lambda_1 + \dots + \lambda_\alpha)^2 - (\lambda_1 + \dots + \lambda_{\alpha-1})^2 \right) = \Psi \left(\lambda_\alpha + 2 \sum_{j=1}^{\alpha-1} \lambda_j \right).$$

(In fact, this is the solution of the full optimization problem with $2^m - 1$ constraints. This fact can be verified, somewhat tediously, by writing the dual of the full problem and directly computing its solution. To shorten the presentation we omit the direct computation and use the above technique.) The optimal value of the cost function, and thus the lower bound on D^* , is given by

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

Applying the definition of Ψ we obtain the desired result. \square

Remark 3.2 (lower bound for all $\varrho \in [0, 1)$). With slight modifications, it is possible to obtain a less tight lower bound valid for all values of ϱ . In the above derivation, the assumption that $\varrho \rightarrow 1^-$ is used only in (3.4). It is possible to use, instead, a lower bound valid for all $\varrho \in [0, 1)$ (see [4]):

$$\bar{d}_\alpha \geq \gamma \sqrt{\frac{|\mathcal{E}|}{\sum_\alpha r_\alpha \bar{N}_\alpha + n/2}},$$

where $\gamma = 2/(3\sqrt{2\pi}) \approx 0.266$. Using this bound we obtain the same linear program as in the proof of Theorem 3.1, with the difference that Ψ is now a function given by

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

Following the procedure in the proof of Theorem 3.1,

$$W_1^* = \frac{\gamma^2 |\mathcal{E}|}{n^2 v^2 (1 - \varrho)^2} \lambda_1 - \frac{n}{2\lambda_1},$$

$$W_\alpha^* = \frac{\gamma^2 |\mathcal{E}|}{n^2 v^2 (1 - \varrho)^2} \left(\lambda_\alpha + 2 \sum_{j=1}^{\alpha-1} \lambda_j \right)$$

for each $\alpha \in \{2, \dots, m\}$. Finally, for every policy P , $D_\alpha(P) \geq W_\alpha^* + \bar{s}_\alpha$, and thus

$$(3.6) \quad D(P) \geq \frac{\gamma^2 |\mathcal{E}|}{n^2 v^2 (1 - \varrho)^2} \sum_{\alpha=1}^m \left(\left(c_\alpha + 2 \sum_{j=\alpha+1}^m c_j \right) \lambda_\alpha \right) - \frac{nc_1}{2\lambda_1} + \sum_{\alpha=1}^m c_\alpha \bar{s}_\alpha$$

for all $\varrho \in [0, 1)$ under the labeling in (2.3).

In the remainder of the paper we design a policy and establish a constant-factor guarantee relative to the heavy-load lower bound.

4. Separate queues policy. In this section we introduce and analyze the separate queues (SQ) policy. We show that this policy is within a factor $2m^2$ 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 $R^{[k]} \subseteq \mathcal{E}$ such that $\{R^{[1]}, \dots, R^{[n]}\}$ forms 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 move to the demand location and provide the on-site service. The SQ policy is described in Algorithm 1. In this algorithm, the probability distribution \mathbf{p} gives a set of parameters which will be used to optimize performance. Without loss of generality, we avoid pathological situations by restricting each p_α to be positive (if $p_\alpha = 0$ for some class α , then the average delay per demand is trivially unbounded).

Figure 4.1 shows an illustrative example of the SQ policy. In the first two frames the vehicle is servicing only class 1 (circle-shaped) demands, whereas in the third frame the vehicle is servicing class 2 (diamond-shaped) demands.

Algorithm 1: SQ Policy

Optimize: algorithm performance over probability distribution
 $\mathbf{p} = [p_1, \dots, p_m]$, where $p_\alpha > 0$ for each $\alpha \in \{1, \dots, m\}$.

- 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 the 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.




FIG. 4.1. A representative simulation of the SQ policy for one vehicle and two priority classes. Circle-shaped demands are high priority, and diamond-shaped are low priority. The vehicle is marked by a chevron-shaped object, and the TSP tour is shown in a solid line. The left figure shows the vehicle computing a tour through class 1 demands. The center figure shows the vehicle partway through the class 1 tour and some newly arrived class 2 demands. The right figure shows the vehicle after completing the class 1 tour and computing a new tour through all class 2 demands.

4.1. 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 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 the beginning of iteration i .

The following straightforward lemma, similar to Lemma 1 in [14], will be essential in deriving our main results.

LEMMA 4.1 (number of outstanding demands). *In heavy-load (i.e., $\rho \rightarrow 1^-$), 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^-).*

Proof. Consider the case where the vehicle moves with infinite velocity (i.e., $v \rightarrow +\infty$); then the system is reduced to an M/G/1 queue (i.e., a queue with exponentially distributed interarrival times, generally distributed service times, and a single server; we refer the reader to [10] for more details). The infinite-velocity system has fewer demands (for every $\alpha \in \{1, \dots, m\}$) waiting in queue. A known result on M/G/1 queues [25] states that, after transients, the total number of demands, as $\rho \rightarrow 1^-$, is very large with high probability. Thus, in the SQ policy, the number of demands in all m classes is very large with high probability. In particular, this implies that, after a transient, the number of demands is very large with high probability at the instances when the vehicle starts a new tour. \square

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

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

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 number of α -demands ($\alpha \in \{1, \dots, m\}$) newly arrived during iteration i ; then, by definition of the SQ policy,

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

By the law of iterated expectation, we have $\mathbb{E}(N_{\alpha,i}^{\text{new}} | TS_j) = (\lambda_\alpha/n)\mathbb{E}(T_i | TS_j)$, where T_i is the length (duration) of the i th iteration. 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

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

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 is large enough, Lemma 4.1 holds, and we can apply (2.1) 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 demands only of type j are serviced), we have

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

where we have

- applied (2.1),

- applied Jensen’s inequality for concave functions in the form $\mathbb{E}[\sqrt{X}] \leq \sqrt{\mathbb{E}[X]}$,
- 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
- used the crucial 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

$$(4.1) \quad \bar{N}_{\alpha,i+1} \leq (1 - p_\alpha)\bar{N}_{\alpha,i} + \sum_{j=1}^m p_j \frac{\lambda_\alpha}{n} \left[\frac{\beta_{\text{TSP}} \sqrt{|\mathcal{E}|}}{\sqrt{nv}} \sqrt{\bar{N}_{j,i} + \bar{N}_{j,i} \bar{s}_j} \right]$$

for each $\alpha \in \{1, \dots, m\}$. The m inequalities above describe 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 4.2 (steady-state queue length). *For every set of initial conditions $\{\bar{N}_{\alpha,0}\}_{\alpha \in \{1, \dots, m\}}$, the trajectories $i \mapsto \bar{N}_{\alpha,i}$, $\alpha \in \{1, \dots, m\}$, resulting from (4.1) satisfy*

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

Proof. Define $q_j := 1 - p_j$, and let $\hat{\lambda}_\alpha$ denote the arrival rate in region $R^{[k]}$. Thus $\hat{\lambda}_\alpha := \lambda_\alpha/n$ for each $\alpha \in \{1, \dots, m\}$. Let $x(i) := (\bar{N}_{1,i}, \bar{N}_{2,i}, \dots, \bar{N}_{m,i}) \in \mathbb{R}^m$, and define two matrices

$$A := \begin{bmatrix} \hat{\lambda}_1 p_1 \bar{s}_1 + q_1 & \hat{\lambda}_1 p_2 \bar{s}_2 & \dots & \hat{\lambda}_1 p_m \bar{s}_m \\ \hat{\lambda}_2 p_1 \bar{s}_1 & \hat{\lambda}_2 p_2 \bar{s}_2 + q_2 & \dots & \hat{\lambda}_2 p_m \bar{s}_m \\ \vdots & & \ddots & \vdots \\ \hat{\lambda}_m p_1 \bar{s}_1 & \hat{\lambda}_m p_2 \bar{s}_2 & \dots & \hat{\lambda}_m p_m \bar{s}_m + q_m \end{bmatrix}$$

and

$$B := \frac{\beta_{\text{TSP}} \sqrt{|\mathcal{E}|}}{\sqrt{nv}} \begin{bmatrix} \hat{\lambda}_1 p_1 & \hat{\lambda}_1 p_2 & \dots & \hat{\lambda}_1 p_m \\ \hat{\lambda}_2 p_1 & \hat{\lambda}_2 p_2 & \dots & \hat{\lambda}_2 p_m \\ \vdots & & \ddots & \vdots \\ \hat{\lambda}_m p_1 & \hat{\lambda}_m p_2 & \dots & \hat{\lambda}_m p_m \end{bmatrix}.$$

Then, letting the relation “ \leq ” in \mathbb{R}^m denote the product order of m copies of \mathbb{R} (in other words, for $v, w \in \mathbb{R}^m$, the relation $v \leq w$ is interpreted componentwise), (4.1) can be written as

$$(4.2) \quad x(i+1) \leq Ax(i) + B \begin{bmatrix} \sqrt{x_1(i)} \\ \sqrt{x_2(i)} \\ \vdots \\ \sqrt{x_m(i)} \end{bmatrix} =: f(x(i)),$$

where $f : \mathbb{R}_{\geq 0}^m \rightarrow \mathbb{R}_{\geq 0}^m$, and $x_j(i)$, $j \in \{1, \dots, m\}$, are the components of vector $x(i)$. We refer to the discrete system in (4.2) as System-X. Next we define two auxiliary

systems, System-Y and System-Z. The initial conditions of these two systems will be set equal to $x(0)$, and we will use their trajectories to bound the trajectories of original system (i.e., System-X). We define System-Y as

$$(4.3) \quad y(i + 1) = f(y(i)).$$

System-Y is, therefore, equal to System-X, with the exception that we replaced the inequality with an equality.

Pick, now, any $\varepsilon > 0$. From Young’s inequality,

$$(4.4) \quad \sqrt{a} \leq \frac{1}{4\varepsilon} + \varepsilon a \quad \text{for all } a \in \mathbb{R}_{\geq 0}.$$

Hence, for $i \mapsto y(i) \in \mathbb{R}_{\geq 0}^m$, (4.3) becomes

$$\begin{aligned} y(i + 1) &\leq Ay(i) + B \left(\frac{1}{4\varepsilon} \mathbf{1}_m + \varepsilon y(i) \right) \\ &= \left(A + \varepsilon B \right) y(i) + \frac{1}{4\varepsilon} B \mathbf{1}_m, \end{aligned}$$

where $\mathbf{1}_m$ is the vector $(1, 1, \dots, 1)^T \in \mathbb{R}^m$. Next, define System-Z as

$$(4.5) \quad z(i + 1) = \left(A + \varepsilon B \right) z(i) + \frac{1}{4\varepsilon} B \mathbf{1}_m =: g(z(i)).$$

The proof now proceeds as follows. First, we show that the initial conditions $x(0) = y(0) = z(0)$ imply that

$$(4.6) \quad x(i) \leq y(i) \leq z(i) \quad \text{for all } i \geq 0.$$

Second, we show that the trajectories of System-Z are bounded; this fact, together with (4.6), implies that also trajectories of System-Y and System-X are bounded. Third, and last, we will compute $\limsup_{i \rightarrow +\infty} y(i)$; this quantity, together with (4.6), will yield the desired result.

Let us consider the first issue. We have $y(1) = f(y(0))$ and $z(1) = g(z(0))$. By definition of System-Y and System-Z, it holds that $z(0) = y(0)$, and thus $g(z(0)) = g(y(0)) \geq f(y(0))$, where the last inequality follows from (4.4) and by definition of f and g . Therefore, we get $y(1) \leq z(1)$. Then, we have $y(2) = f(y(1))$ and $z(2) = g(z(1))$. Since $z(1), y(1) \in \mathbb{R}_{\geq 0}^m$ and the elements in matrices A and B are all nonnegative, $y(1) \leq z(1)$ implies $g(y(1)) \leq g(z(1))$. Using similar arguments, we can write $z(2) \geq g(y(1)) \geq f(y(1)) = y(2)$; therefore, we get $y(2) \leq z(2)$. Then, it is immediate by induction that $y(i) \leq z(i)$ for all $i \geq 0$.

Similarly, by definition of System-Y, it holds that $x(0) = y(0)$, and thus $x(1) \leq f(x(0)) = f(y(0)) = y(1)$. Then, we get $x(1) \leq y(1)$. Since $x(1), y(1) \in \mathbb{R}_{\geq 0}^m$ and the elements in matrices A and B are nonnegative, and by the monotonicity of $\sqrt{\cdot}$, $x(1) \leq y(1)$ implies $f(x(1)) \leq f(y(1))$. Therefore, we can write $x(2) \leq f(x(1)) \leq f(y(1)) = y(2)$; thus, we get $x(2) \leq y(2)$. Then, it is immediate to show by induction that $x(i) \leq y(i)$ for all $i \geq 0$, and (4.6) holds.

We now turn our attention to the second issue, namely boundedness of trajectories for System-Z (in (4.5)). Notice that System-Z is a discrete-time linear system. The eigenvalues of A are characterized in the following lemma.

LEMMA 4.3. *The eigenvalues of A are real and have magnitude strictly less than 1 (i.e., A is a stable matrix).*

Proof. Let $w \in \mathbb{C}^m$ be an eigenvector of A , and let $\mu \in \mathbb{C}$ be the corresponding eigenvalue. Then we have $Aw = \mu w$. Define $r := (p_1 \bar{s}_1, p_2 \bar{s}_2, \dots, p_m \bar{s}_m)$. Then the m eigenvalue equations are

$$(4.7) \quad \hat{\lambda}_j w \cdot r + q_j w_j = \mu w_j, \quad j \in \{1, \dots, m\},$$

where $w \cdot r$ is the scalar product of vectors w and r , and w_j is the j th component of w .

There are two possible cases. The first case is that $w \cdot r = 0$. (Note that since each $p_\alpha > 0$, this case can occur only if $\bar{s}_\alpha = 0$ for some $\alpha \in \{1, \dots, m\}$.) In this case, (4.7) becomes $q_j w_j = \mu w_j$ for all j . Since $w \neq 0$, there exists j^* such that $w_{j^*} \neq 0$; thus, we have $\mu = q_{j^*}$. Since $q_{j^*} \in \mathbb{R}$ and $0 < q_{j^*} < 1$, we have that μ is real and $|\mu| < 1$.

Assume, now, that $w \cdot r \neq 0$. This implies that $\mu \neq q_j$ and $w_j \neq 0$ for all j ; thus we can write, for all j ,

$$(4.8) \quad w_j = \frac{\hat{\lambda}_j}{\mu - q_j} w \cdot r,$$

and hence

$$w_j = \frac{\hat{\lambda}_j \mu - q_1}{\hat{\lambda}_1 \mu - q_j} w_1.$$

Therefore, (4.8) can be rewritten as

$$(4.9) \quad \sum_{j=1}^m \frac{r_j \hat{\lambda}_j}{\mu - q_j} = 1.$$

Equation (4.9) implies that the eigenvalues are real. To see this, write $\mu = a + ib$, where i is the imaginary unit: then

$$\sum_{j=1}^m \frac{r_j \hat{\lambda}_j}{a + ib - q_j} = \sum_{j=1}^m \frac{r_j \hat{\lambda}_j [(a - q_j) - ib]}{(a - q_j)^2 + b^2}.$$

Thus (4.9) implies

$$b \underbrace{\sum_{j=1}^m \frac{r_j \hat{\lambda}_j}{(a - q_j)^2 + b^2}}_{>0} = 0,$$

that is, $b = 0$. Equation (4.9) also implies that the eigenvalues (that are real) have magnitude strictly less than 1. Indeed, assume, by contradiction, that $\mu \geq 1$. Then we have $\mu - q_j \geq 1 - q_j > 0$ (recall that the eigenvalues are real and $0 < q_j < 1$) and we can write

$$\sum_{j=1}^m \frac{r_j \hat{\lambda}_j}{\mu - q_j} \leq \sum_{j=1}^m \frac{r_j \hat{\lambda}_j}{1 - q_j} = \sum_{j=1}^m \bar{s}_j \hat{\lambda}_j = \varrho < 1,$$

which is a contradiction. Assume, again by contradiction, that $\mu \leq -1$. In this case we trivially get another contradiction $\sum_{j=1}^m r_j \hat{\lambda}_j / (\mu - q_j) < 0$, since $\mu - q_j < 0$. \square

Hence, $A \in \mathbb{R}^{m \times m}$ has eigenvalues strictly inside the unit disk, and since the eigenvalues of a matrix depend continuously on the matrix entries, there exists a sufficiently small $\varepsilon > 0$ such that the matrix $A + \varepsilon B$ has eigenvalues strictly inside the unit disk. Accordingly, each solution $i \mapsto z(i) \in \mathbb{R}_{\geq 0}^m$ of System-Z converges exponentially fast to the unique equilibrium point

$$(4.10) \quad z^* = \left(I_m - A - \varepsilon B \right)^{-1} \frac{1}{4\varepsilon} B \mathbf{1}_m.$$

Combining (4.6) with the previous statement, we see that the solutions $i \mapsto x(i)$ and $i \mapsto y(i)$ are bounded. Thus

$$(4.11) \quad \limsup_{i \rightarrow +\infty} x(i) \leq \limsup_{i \rightarrow +\infty} y(i) < +\infty.$$

Finally, we turn our attention to the third issue, namely the computation of $y := \limsup_{i \rightarrow +\infty} y(i)$. Taking the lim sup of the left- and right-hand sides of (4.3), and noting that

$$\limsup_{i \rightarrow +\infty} \sqrt{y_\alpha(i)} = \sqrt{\limsup_{i \rightarrow +\infty} y_\alpha(i)} \quad \text{for } \alpha \in \{1, 2, \dots, m\},$$

since $x \mapsto \sqrt{x}$ is continuous and strictly monotone increasing on $\mathbb{R}_{>0}$, we obtain that

$$y_\alpha = (1 - p_\alpha)y_\alpha + \hat{\lambda}_\alpha \sum_{j=1}^m p_j \left(\frac{\beta_{\text{TSP}} \sqrt{|\mathcal{E}|}}{\sqrt{nv}} \sqrt{y_j} + \bar{s}_j y_j \right).$$

Rearranging we obtain

$$(4.12) \quad p_\alpha y_\alpha = \hat{\lambda}_\alpha \sum_{j=1}^m p_j \left(\frac{\beta_{\text{TSP}} \sqrt{|\mathcal{E}|}}{\sqrt{nv}} \sqrt{y_j} + \bar{s}_j y_j \right).$$

Dividing $p_\alpha y_\alpha$ by $p_1 y_1$ (recall that by Algorithm 1 each p_α is positive) we obtain

$$(4.13) \quad y_\alpha = \frac{\hat{\lambda}_\alpha p_1}{\hat{\lambda}_1 p_\alpha} y_1.$$

Combining (4.12) and (4.13), we obtain

$$p_1 y_1 = \varrho p_1 y_1 + \frac{\beta_{\text{TSP}} \sqrt{|\mathcal{E}|}}{\sqrt{nv}} \sqrt{p_1 \hat{\lambda}_1 y_1} \sum_{j=1}^m \sqrt{\hat{\lambda}_j p_j}.$$

Thus, recalling that $\hat{\lambda}_\alpha = \lambda_\alpha/n$, we obtain

$$y_\alpha = \frac{\beta_{\text{TSP}}^2 |\mathcal{E}|}{n^3 v^2 (1 - \varrho)^2} \frac{\lambda_\alpha}{p_\alpha} \left(\sum_{j=1}^m \sqrt{\lambda_j p_j} \right)^2.$$

Noting that, from (4.11), $\limsup_{i \rightarrow +\infty} \bar{N}_{\alpha,i} \leq y_\alpha$, we obtain the desired result. \square

4.2. Delay of the SQ policy in heavy-load. From Theorem 4.2, and using Little's law, the delay of α -demands satisfies

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

where we have neglected \bar{s}_α , since as $\rho \rightarrow 1^-$ the constant \bar{s}_α becomes negligible compared to the average delay, which scales as $(1-\rho)^{-2}$.

Thus, the delay of the SQ policy satisfies

$$(4.14) \quad D(\text{SQ}) \leq \frac{\beta_{\text{TSP}}^2 |\mathcal{E}|}{n^2 v^2 (1-\rho)^2} \sum_{\alpha=1}^m \frac{c_\alpha}{p_\alpha} \left(\sum_{i=1}^m \sqrt{\lambda_i p_i} \right)^2 \quad \text{as } \rho \rightarrow 1^-.$$

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

THEOREM 4.4 (SQ policy performance). *As $\rho \rightarrow 1^-$, the delay of the SQ policy is within a factor $2m^2$ of the optimal delay. This factor is 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

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

Defining $\Psi := \beta_{\text{TSP}}^2 |\mathcal{E}| / (n^2 v^2 (1-\rho)^2)$, (4.14) can be written as

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

Next, the lower bound in (3.1) is

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

Thus, comparing the upper and lower bounds,

$$(4.15) \quad \frac{D(\text{SQ})}{D^*} \leq 2m \frac{\left(\sum_{i=1}^m \sqrt{c_i \lambda_i} \right)^2}{\sum_{i=1}^m (c_i \lambda_i)}.$$

Letting $x_i := \sqrt{c_i \lambda_i}$, and $\mathbf{x} := [x_1, \dots, x_m]$, the numerator of the fraction in (4.15) 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,

$$\frac{D(\text{SQ})}{D^*} \leq 2m \left(\frac{\|\mathbf{x}\|_1}{\|\mathbf{x}\|_2} \right)^2 \leq 2m^2 \quad \text{as } \rho \rightarrow 1^-,$$

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

Remark 4.5 (relation to RP policy in [21]). 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 [21]. 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.

4.3. SQ policy with queue merging. In this section we propose a modification of the SQ policy based on *queue merging*. Queue merging is guaranteed to never increase the upper bound on the expected delay, and in certain instances it significantly decreases the upper bound. The modification can be used when we have a modest number of classes (fewer than, say, 20), which encompasses most applications of interest.

To motivate the modification, consider the case when all classes have equal priority (i.e., $c_1/\lambda_1 = \dots = c_m/\lambda_m$), and we use the probability assignment $p_\alpha = c_\alpha$ for each class α . Then, the upper bound for the SQ policy in (4.14) becomes

$$\Psi m(\lambda_1 + \dots + \lambda_m),$$

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

On the other hand, if we ignore priorities, merge the m classes into a single class, and run the SQ policy on the merged class (i.e., at each iteration, service all outstanding demands in \mathcal{E} via the TSP tour), then the upper bound becomes

$$\Psi(\lambda_1 + \dots + \lambda_m).$$

Thus, there is a factor of m separating the two upper bounds. This is due to the fact that the basic SQ policy services each of the m classes separately, even when they have the same priority.

The above discussion motivates the addition of queue merging to the SQ policy. We define a *merge configuration* to be a partition of m classes $\{1, \dots, m\}$ into ℓ sets C_1, \dots, C_ℓ , where $\ell \in \{1, \dots, m\}$. The upper bound for a merge configuration $\{C_1, \dots, C_\ell\}$ is

$$(4.16) \quad \Psi \ell \left(\sum_i^\ell \sqrt{\sum_{\alpha \in C_i} c_\alpha \sum_{\beta \in C_i} \lambda_\beta} \right)^2.$$

The SQ policy with merging can be summarized as follows.

SQ with Merging Policy

- 1 Find the merge configuration $\{C_1, \dots, C_\ell\}$ which minimizes (4.16).
 - 2 Run the SQ policy on ℓ classes, where class i has arrival rate $\sum_{\alpha \in C_i} \lambda_\alpha$ and convex combination coefficient $\sum_{\alpha \in C_i} c_\alpha$.
-

Now, to minimize (4.16) in step 1 of the SQ with merging policy, one must search over all possible partitions of a set of m elements. The number of partitions is given by the Bell number B_m , which is defined recursively as $B_m = \sum_{k=0}^{m-1} B_k \binom{m-1}{k}$. Thus, the search becomes infeasible for more than 10 classes.

If the search space is too large, then one can limit the search to all partitions such that if $i < j$, then each class in C_i has higher priority than all classes in C_j . This is the set of partitions in which only adjacent classes are merged. For m classes, there are 2^{m-1} such merge configurations, which is significantly less than the Bell number B_m but still infeasible for more than, say, 20 classes.

4.4. The “tube” heuristic for improving performance. We now introduce a simple heuristic improvement for the SQ policy that can be used for implementation. The heuristic improvement is as follows:




FIG. 4.2. The tube heuristic for two classes of demands with $c = 0.8$, $\lambda_2 = 6\lambda_1$, and several different load factors ρ . The delay at $\epsilon = 0$ corresponds to the basic SQ policy.

The tube heuristic. When following the tour in step 10 of the SQ policy, service all newly arrived demands that lie within distance $\epsilon > 0$ of the tour.

The idea behind the heuristic is to utilize the fact that some newly arrived demands will be “close” to the demands in the current service batch and thus can be serviced with minimal additional travel cost. Analysis of the tube heuristic is complicated by the fact that it introduces correlation between demand locations. A similar difficulty arises when attempting to analyze the nearest neighbor policy [3]. However, we can demonstrate the effectiveness of this heuristic through simulations.

The parameter ϵ should be chosen such that the total tour length is not increased by more than, say, 10%. A rough calculation shows that the area of the “tube” of width 2ϵ centered around a tour that passes through the $\text{card}(Q)$ demands in Q has area upper bounded by $2\epsilon\beta_{\text{TSP}}\sqrt{\text{card}(Q)|\mathcal{E}|}$. While following the tour, a vehicle will deviate to service no more than $2\epsilon\beta_{\text{TSP}}\sqrt{\text{card}(Q)/|\mathcal{E}|}(\bar{N}_1 + \dots + \bar{N}_m)$ demands. Finally, since the vehicle will have to travel no more than 2ϵ to service each demand in the “tube,” we see that ϵ should scale as

$$\epsilon \sim \sqrt{\frac{f|\mathcal{E}|}{\bar{N}_1 + \dots + \bar{N}_m}},$$

where \bar{N}_α is the expected number of α -demands in the environment, and f is the fractional increase in tour length (e.g., 10%).

Figure 4.2 shows numerical results for the tube heuristic for a single unit speed vehicle in a square environment with side length 50. The simulation is performed for two classes of demands with $c = 0.8$, $\lambda_2 = 6\lambda_1$, and several different load factors ρ . Each experimental data point represents the average of the steady-state delay of 10 runs, where each run consists of 200 iterations of the SQ policy. To ensure convergence to steady state and avoid effects due to the transient response, only the last 50 iterations in each run are used to calculate the delay. The basic policy is shown in the leftmost data points (i.e., $\epsilon = 0$). Figure 4.2 demonstrates that as the load factor increases, the value of ϵ should be chosen smaller in order to achieve the best performance. Table 4.1 shows the improvement in expected delay when using the tube heuristic. For the load factors considered, the heuristic decreases the delay by a

TABLE 4.1

A comparison between the expected delay of the basic SQ policy and the SQ policy with the tube heuristic. The values in brackets give the standard deviation of the corresponding table entry.

Load factor ρ	Delay	Best ϵ	Delay with best ϵ	Heuristic improvement
0.14	358 (34)	5	183 (11)	0.51 (0.16)
0.28	496 (61)	4	244 (25)	0.49 (0.23)
0.42	774 (78)	3	384 (26)	0.50 (0.17)
0.56	1330 (84)	2	706 (52)	0.53 (0.14)
0.70	3380 (357)	1	1770 (121)	0.52 (0.17)

factor of approximately 2. One should note that the heuristic is difficult to accurately simulate for high load factors. This is due to the additional computations required to determine if a newly arrived demand lies within a distance ϵ of the current tour. A more sophisticated implementation of the tube heuristic is to define an ϵ_α for each $\alpha \in \{1, \dots, m\}$, where the magnitude of ϵ_α is proportional to its priority and thus proportional to the probability p_α .

5. Simulations and discussion. In this section we discuss, through the use of simulations, the performance of the SQ policy with the probability assignment $p_\alpha := c_\alpha$ for each $\alpha \in \{1, \dots, m\}$. In particular, we study (i) the tightness of the upper bound in (4.14), (ii) conditions for which the gap between the lower bound in (3.1) and the upper bound in (4.14) is maximized, (iii) the suboptimality of the probability assignment $p_\alpha = c_\alpha$, and (iv) the difference in performance between the SQ policy and a policy that merges all classes 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.

5.1. Tightness of the upper bound. We consider one vehicle, four classes of demands, and several values of the load factor ρ . For each value of ρ we perform 100 runs. In each run we uniformly randomly generate arrival rates $\lambda_1, \dots, \lambda_m$, convex combination coefficients c_1, \dots, c_m , and on-site service times $\bar{s}_1, \dots, \bar{s}_m$ and normalize the values such that the constraints $\sum_{\alpha=1}^m \lambda_\alpha \bar{s}_\alpha = \rho$ and $\sum_{\alpha=1}^m c_\alpha = 1$ are satisfied. In each run we iterate the SQ policy 4000 times and compute the steady-state expected delay by considering the number of demands in the last 1000 iterations. For each value of ρ we compute the ratio χ between the expected delay and the theoretical upper bound in (4.14). Table 5.1 reports the ratio, its standard deviation, and its minimum and maximum values for each ρ value. One can see that the upper bound provides a reasonable approximation for load factors as low as $\rho = 0.75$.

TABLE 5.1

Ratio χ between experimental results and the upper bound for various values of ρ .

Load factor (ρ)	$\mathbb{E}[\chi]$	σ_χ	$\max \chi$	$\min \chi$
0.75	0.803	0.092	1.093	0.354
0.8	0.778	0.108	0.943	0.256
0.85	0.773	0.111	1.150	0.417
0.9	0.733	0.159	1.162	0.203
0.95	0.716	0.131	0.890	0.257

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




FIG. 5.1. Experimental results for the SQ policy in worst-case conditions plotted on a log-log scale; $\rho = 0.85$ and $\lambda_1 = 1$.

5.2. Maximum deviation from lower bound. In Theorem 4.4 we showed that the SQ policy performs within a factor of $2m^2$ of the lower bound for all initial conditions. The ratio between the upper bound (4.14) and the lower bound in (3.1) can be made arbitrarily close to $2m^2$ by choosing $\lambda_1 \ll \lambda_2 \ll \dots \ll \lambda_m$ and $c_1 \gg c_2 \gg \dots \gg c_m$, with $\lambda_\alpha c_\alpha = a$, for each $\alpha \in \{1, \dots, m\}$ and for some positive constant a . In these “unfavorable conditions,” the upper bound is equal to Bm^3a and the lower bound is approximately $Bma/2$.

It is also of interest to consider the actual deviation of the experimental performance from the lower bound in the unfavorable conditions described above. We simulated the SQ policy for $\rho = 0.85$ and for several values of m , with parameter values of $\lambda_m = b\lambda_{m-1} = b^2\lambda_{m-2} = \dots = b^{m-1}\lambda_1$ and $c_1 = bc_2 = \dots = b^{m-1}c_m$, where $b = 2$. Figure 5.1 (plotted on a log-log scale) shows that the ratio between the actual performance and the lower bound (averaged over 10 simulation runs) increases as m^η , where $\eta \approx 1.25$ according to a least square fit. The figure also shows that the ratio between the analytic upper bound and the lower bound increases as m^η , where $\eta \approx 1.61$ according to a least square fit. For completeness, the figure also shows the worst-case ratio between the upper bound and lower bound, which increases as m^2 . These experimental results suggest that the upper bound is somewhat conservative.

5.3. Suboptimality of the approximate probability assignment. To prove Theorem 4.4 we used the probability assignment

$$(5.1) \quad p_\alpha := c_\alpha \quad \text{for each } \alpha \in \{1, \dots, m\}.$$

However, one would like to select $[p_1, \dots, p_m] =: \mathbf{p}$, which minimizes the right-hand side of (4.14). The minimization of the right-hand side of (4.14) is a constrained multivariable nonlinear optimization problem over \mathbf{p} , that is, in m dimensions. Thus, for a general m class problem, solving the optimization problem is difficult. However, for two classes of demands the optimization is over a single variable p_1 (with the constraint that $p_2 = 1 - p_1$), and it can be readily solved. A comparison of the optimized upper bound, denoted upbd_{opt} , with the upper bound obtained using the probability assignment in (5.1), denoted upbd_c , is shown in Figure 5.2. In this figure the ratio of upper bounds is bounded by two.

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




FIG. 5.2. The ratios $upbd_c/upbd_{opt}$ for two classes of demands.

TABLE 5.2

Ratio of the upper bound with $p_\alpha = c_\alpha$ for each $\alpha \in \{1, \dots, m\}$ and the upper bound with a locally optimized probability assignment.

Number of classes (m)	$upbd_c/upbd_{local\ opt}$	Max. % variation in ratio
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

$upbd_{local\ opt}$ obtained from the five locally optimized probability assignments. We also record the maximum variation in the five locally optimized upper bounds. This is summarized in Table 5.2. 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 (5.1) seems to perform within a factor of two of the optimized assignment, and 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.

5.4. The complete merge policy. As described in section 4.3, a naive policy for our problem is to ignore priorities, merge all classes into a single class, and repeatedly form TSP tours through all outstanding demands. We call this policy the complete merge (CM) policy. In this section we briefly verify by simulation that the performance of the CM policy can be very poor when compared to that of the SQ policy. In addition, the poor performance occurs under conditions of interest for most applications—when low-priority demands arrive much more frequently than high-priority demands.

To upper bound the performance of the CM policy, define the total arrival rate $\Lambda := \sum_{\alpha=1}^m \lambda_\alpha$ and total mean on-site service $\bar{S} := \sum_{\alpha=1}^m \bar{s}_\alpha$. Using the upper bounds in [3], we immediately obtain that $D(\text{CM}) \leq \frac{\beta_{\text{TSP}}^2 \epsilon |\Lambda|}{n^2 v^2 (1-\theta)^2}$. Thus, the ratio of upper bounds can be made arbitrarily large by choosing $\lambda_1 \ll \lambda_2 \ll \dots \ll \lambda_m$ and $c_1 \gg$




FIG. 5.3. Ratio of experimental delays between the CM policy and the SQ policy as a function of λ_2 , with $m = 2$, $\lambda_1 = 1$, $c = 0.995$, and $\rho = 0.9$.

$c_2 \gg \dots \gg c_m$. This suggests that the ratio between the delay of the CM policy and that of the SQ policy, $D(\text{CM})/D(\text{SQ})$, can be made very large. Figure 5.3 shows the experimentally obtained ratio between the delay of the CM policy and that of the SQ policy (averaged over 10 simulation runs) and verifies that the above choice of arrival rates and convex combination coefficients results in large performance ratios.

6. Conclusions. In this paper we introduced a DVR problem with priority classes. We captured the priority levels of classes by writing the system delay as a convex combination of the delay of each class. We determined a lower bound on the achievable values of the convex combination of the class delays. We then presented the SQ policy and showed that it performs within a constant factor of the lower bound, which depends only on the number of classes. We believe that it may be possible to improve the lower bound and remove, or reduce, the constant factor's dependence on the number of classes. For future work we are interested in combining the aspects of multiclass vehicle routing with problems in which demands require teams of vehicles for their service. We are also interested in extending our results to the case of nonuniform demand densities (possibly class dependent) and to impatient demands that disappear if they are not serviced within a certain time window.

REFERENCES

- [1] R. W. BEARD, T. W. MCLAIN, M. A. GOODRICH, AND E. P. ANDERSON, *Coordinated target assignment and intercept for unmanned air vehicles*, IEEE Trans. Robotics and Automation, 18 (2002), pp. 911–922.
- [2] D. BERTSIMAS, I. C. PASCHALIDIS, AND J. N. TSITSIKLIS, *Optimization of multiclass queueing networks: Polyhedral and nonlinear characterizations of achievable performance*, Ann. Appl. Probab., 4 (1994), pp. 43–75.
- [3] D. J. BERTSIMAS AND G. J. VAN RYZIN, *A stochastic and dynamic vehicle routing problem in the Euclidean plane*, Oper. Res., 39 (1991), pp. 601–615.
- [4] D. J. BERTSIMAS AND G. J. VAN RYZIN, *Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles*, Oper. Res., 41 (1993), pp. 60–76.
- [5] A. BLUM, P. CHALASANI, D. COPPERSMITH, B. PULLEYBLANK, P. RAGHAVAN, AND M. SUDAN, *The minimum latency problem*, in Proceedings of the ACM Symposium on the Theory of Computing, Montreal, Canada, 1994, pp. 163–171.
- [6] A. BLUM, S. CHAWLA, D. R. KARGER, T. LANE, A. MEYERSON, AND M. MINKOFF, *Approximation algorithms for orienteering and discounted-reward TSP*, SIAM J. Comput., 37 (2007), pp. 653–670.
- [7] 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 Trans. Automat. Control, to appear.

- [8] E. G. COFFMAN, JR., AND I. MITRANI, *A characterization of waiting time performance realizable by single-server queues*, Oper. Res., 28 (1980), pp. 810–821.
- [9] E. FRAZZOLI AND F. BULLO, *Decentralized algorithms for vehicle routing in a stochastic time-varying environment*, in Proceedings of the IEEE Conference on Decision and Control, Paradise Island, Bahamas, 2004, pp. 3357–3363.
- [10] L. KLEINROCK, *Queueing Systems. Volume II: Computer Applications*, John Wiley, New York, 1976.
- [11] A. LARSEN, *The Dynamic Vehicle Routing Problem*, Ph.D. thesis, Technical University of Denmark, Lyngby, Denmark, 2000.
- [12] R. C. LARSON AND A. R. ODoni, *Urban Operations Research*, Prentice–Hall, Englewood Cliffs, NJ, 1981.
- [13] D. G. LUENBERGER, *Linear and Nonlinear Programming*, 2nd ed., Addison–Wesley, Reading, MA, 1984.
- [14] J. D. PAPASTAVROU, *A stochastic and dynamic routing policy using branching processes with state dependent immigration*, European J. Oper. Res., 95 (1996), pp. 167–177.
- [15] M. PAVONE, N. BISNIK, E. FRAZZOLI, AND V. ISLER, *A stochastic and dynamic vehicle routing problem with time windows and customer impatience*, ACM/Springer J. Mobile Networks and Applications, 14 (2009), pp. 350–364.
- [16] M. PAVONE, E. FRAZZOLI, AND F. BULLO, *Distributed and adaptive algorithms for vehicle routing in a stochastic and dynamic environment*, IEEE Trans. Automat. Control, to appear; available online from <http://arxiv.org/abs/0903.3624>.
- [17] M. PAVONE, S. L. SMITH, F. BULLO, AND E. FRAZZOLI, *Dynamic multi-vehicle routing with multiple classes of demands*, in Proceedings of the American Control Conference, St. Louis, MO, 2009, pp. 604–609.
- [18] A. G. PERCUS AND O. C. MARTIN, *Finite size and dimensional dependence of the Euclidean traveling salesman problem*, Phys. Rev. Lett., 76 (1996), pp. 1188–1191.
- [19] H. N. PSARAFTIS, *Dynamic vehicle routing problems*, in Vehicle Routing: Methods and Studies, B. Golden and A. Assad, eds., North–Holland, Amsterdam, 1988, pp. 223–248.
- [20] S. L. SMITH AND F. BULLO, *The dynamic team forming problem: Throughput and delay for unbiased policies*, Systems Control Lett., 58 (2009), pp. 709–715.
- [21] S. L. SMITH, M. PAVONE, F. BULLO, AND E. FRAZZOLI, *Dynamic traveling repairperson with priority demands*, in Proceedings of the IEEE Conference on Decision and Control, Cancún, México, 2008, pp. 1206–1211.
- [22] M. Z. SPIVEY AND W. B. POWELL, *The dynamic assignment problem*, Transportation Sci., 38 (2004), pp. 399–419.
- [23] J. M. STEELE, *Probabilistic and worst case analyses of classical problems of combinatorial optimization in Euclidean space*, Math. Oper. Res., 15 (1990), pp. 749–770.
- [24] H. A. WAISANEN, D. SHAH, AND M. A. DAHLEH, *A dynamic pickup and delivery problem in mobile networks under information constraints*, IEEE Trans. Automat. Control, 53 (2008), pp. 1419–1433.
- [25] R. W. WOLFF, *Stochastic Modeling and the Theory of Queues*, Prentice–Hall, Englewood Cliffs, NJ, 1989.
- [26] H. XU, *Optimal Policies for Stochastic and Dynamic Vehicle Routing Problems*, Ph.D. thesis, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA, 1995.