

STOCHASTIC AND DYNAMIC VEHICLE ROUTING WITH GENERAL DEMAND AND INTERARRIVAL TIME DISTRIBUTIONS

DIMITRIS J. BERTSIMAS,* *Sloan School of Management, MIT*

GARRETT VAN RYZIN,** *Columbia University*

Abstract

We analyze a class of stochastic and dynamic vehicle routing problems in which demands arrive randomly over time and the objective is minimizing waiting time. In our previous work ([6], [7]), we analyzed this problem for the case of uniformly distributed demand locations and Poisson arrivals. In this paper, using quite different techniques, we are able to extend our results to the more realistic case where demand locations have an arbitrary continuous distribution and arrivals follow only a general renewal process. Further, we improve significantly the best known lower bounds for this class of problems and construct policies that are provably within a small constant factor relative to the optimal solution. We show that the leading behavior of the optimal system time has a particularly simple form that offers important structural insight into the behavior of the system. Moreover, by distinguishing two classes of policies our analysis shows an interesting dependence of the system performance on the demand distribution.

GENERAL DEMAND DISTRIBUTION; GENERAL INTERARRIVAL TIMES; BOUNDS; HEURISTICS

AMS 1991 SUBJECT CLASSIFICATION: PRIMARY 60D05

SECONDARY 60K25, 60K30

1. Introduction

Dynamic vehicle routing problems occur when one has to visit customers (*demands*) that arrive sequentially over time. The objective is to schedule these visits in a way that is economical, yet also provides an acceptable service level (wait for delivery or service). Because future demand is often uncertain, it is natural to view the sequence of arrivals in dynamic vehicle routing problems as a stochastic process. That is, at time t we know the location and age of all demands that have arrived before time t , but we have only a probabilistic characterization of future demand locations and arrival times. In addition, if the planning horizon is sufficiently long (as in the day-to-day operation of a distribution facility), we may

Received 27 April 1992.

* Postal address: Sloan School of Management, MIT, Room E53-359, Cambridge, MA 02139, USA.

** Postal address: Graduate School of Business, Columbia University, New York, NY 10027, USA.

The research of D. Bertsimas was partially supported by a Presidential Young Investigator Award DDM-9158118 with matching funds from Draper Laboratory. The research of both authors was partially supported by the National Science Foundation under grant DDM-9014751 and by grants from Draper Laboratory and the UPS Foundation.

view the problem as that of finding a *stationary policy* for scheduling vehicles that minimizes a *time average* cost over an *infinite horizon*. In such a setting, the economical delivery objective corresponds to minimizing the average distance traveled per demand served, while the service-level objective corresponds to minimizing the average wait for delivery or service.

There are many practical settings in which such problems arise. Any distribution system that receives orders in real time and makes deliveries based on these orders (courier services, deliveries of flowers, pizza, etc.) is a clear candidate. Other applications include scheduling repair crews to service geographically dispersed failures. Examples of this type include repairing electric utility networks, contract maintenance of customer premise equipment by computer and telecom equipment vendors, and road and highway maintenance. Distribution of finished goods from factories to retailers is a third important application area. Arrivals in this context correspond to the completion of a unit at the factory that is designated for a particular retailer. In this case, the average waiting time represents a service level and, for a fixed production rate, a measure of the average inventory held in the distribution channel. A specific example is the delivery of automobiles from assembly plants to local dealers. We refer the reader to the excellent general discussion in [17] and to the more specific discussions in [6] and [7] for more background on dynamic vehicle routing problems.

1.1. *The dynamic traveling repairman problem (DTRP)*. In [6] and [7] we analyzed a version of a dynamic vehicle routing problem, which we call the dynamic traveling repairman problem (DTRP). Demands arrive according to a Poisson process to a Euclidean service region, and their locations are independent and uniformly distributed throughout the service region. (These locations are random only in the sense mentioned above; namely that at time t we know with *certainty* the locations of demands that arrived before time t , but future demand locations form an i.i.d. sequence.) At each location, the vehicle serving the demand must spend some amount of time in on-site service. We assume this time is a generally distributed, i.i.d. random variable for all demands that is realized only when service is completed. The objective is to find a stationary policy μ for routing one or more vehicles that travel at constant velocity so that the average system time (wait for completion of service), T_μ , is minimized; that is, we seek a policy that *maximizes* the *level of service* provided by the fleet of vehicles.

For the DTRP, we found policies that were provably optimal in light traffic and policies that had system times whose ratio to the optimum system time, T^* , were provably within a constant factor in heavy traffic. The best of these later policies, the so-called modified TSP policy [7], has a guarantee of $T_\mu/T^* \leq 3.6$. In [7] we also extended the model to the case where vehicles have an upper bound q on the number of demands they can serve before having to return to a given depot location.

1.2. *Overview and contribution of this paper*. In this paper, we extend the analysis of the DTRP to the case where demand locations are distributed to an

arbitrary, continuous density defined over the service region and arrivals form a general renewal process. (See Section 2 for formal definitions and assumptions.) These extensions are important for both practical and theoretical reasons. In practice, one rarely encounters uniformly distributed locations; therefore, to actually use such models, an understanding of the problem cost and solution structure under general distributions is essential. Further, while the Poisson arrival assumption may be appropriate for certain repair systems, other applications (e.g. finished goods distribution from assembly plants) have either more or less variable interarrival times: therefore, it is important to understand to what extent the results hold for non-Poisson arrivals.

From a theoretical perspective, these generalizations are also quite challenging. In our previous analysis, uniformity and PASTA [20] were heavily exploited. We are therefore forced to develop new proof techniques that do not rely on these properties. These new techniques are introduced and developed in Section 3. This new approach has the added benefit of strengthening the lower bound on T^* by a factor of 2. Thus, we are able to improve our heavy traffic guarantee for the modified TSP policy to $T_\mu/T^* \leq 1.8$, which brings the guarantee into a more practical range. Indeed, we conjecture that the modified TSP policy is asymptotically optimal in heavy traffic.

These extensions also reveal interesting structural properties of the problem that are not apparent in the uniform case. Specifically, it turns out that we need to distinguish between policies that provide the same level of service (i.e. mean waiting time) for all locations, which we call *spatially unbiased* policies, and those which may produce waiting times that vary with location, which we call *spatially biased* policies. We fully characterize these two behaviors and construct policies that are provably good for both in Section 4. In Section 5, we compare these two behaviors and show that there is always an increase in the optimal system time if one requires spatially unbiased service. Further, this gap widens as the distribution of locations becomes ‘less uniform’.

We re-examine the lower bounds in Section 6 and show that they are as tight as possible given the information used in their derivation. Thus, one can improve this bound only by exploiting more of the problem’s vehicle routing structure. Finally, in Section 7 we briefly mention some other extensions to the case of capacitated vehicle and higher-dimensional spaces. In Section 8, we give our conclusions.

2. Notation and definitions

2.1. Problem definition. A total of m identical vehicles operates in a service region $\mathcal{A} \subset \mathbb{R}^2$. Vehicles travel at a constant, finite velocity v . Demands arrive to \mathcal{A} according to a stationary renewal process with intensity λ . A demand’s location becomes known at its arrival epoch, and it stays at its location until a vehicle visits it and completes its on-site service. We make the following assumptions in the model:

1. $\mathcal{A} \subset \mathbb{R}^2$ is closed and bounded.

- 2. Interarrival times have finite variance σ_a^2 and Laplace transform $A^*(s)$.
- 3. On-site service times $\{s_i; i \geq 1\}$ are generally distributed i.i.d. random variables with finite first and second moments denoted by \bar{s} and \bar{s}^2 respectively.
- 4. Demand locations $\{X_i; i \geq 1\}$ are i.i.d. and distributed according to a continuous density $f(x)$ which satisfies

$$P\{X_i \in \mathcal{S}\} = \int_{\mathcal{S}} f(x) dx \quad \forall \mathcal{S} \subseteq \mathcal{A} \quad \text{and} \quad \int_{\mathcal{A}} f(x) dx = 1.$$

- 5. The density $f(x)$ is K -Lipschitz,

$$(1) \quad |f(x) - f(y)| \leq K \|x - y\| \quad \forall x, y \in \mathcal{A},$$

and bounded above and below,

$$(2) \quad 0 < \underline{f} \leq f(x) \leq \bar{f} < \infty, \quad \forall x \in \mathcal{A}.$$

A policy for routing the vehicles is said to be *stable* if the expected number of demands in the system is bounded uniformly for all times t . If a policy is stable, $\rho \equiv \lambda \bar{s} / m$ is the fraction of total vehicle time spent in on-site service. The term *heavy traffic* denotes the condition $\rho \rightarrow 1$. We let W denote the mean waiting time and $T = W + \bar{s}$ denote the mean system time (wait in queue plus on-site service) of a randomly selected demand.

Let \mathcal{M} denote the subset of all stable, stationary policies μ with the following properties:

- P1. Decisions are taken only at service completion epochs except for vehicles that are at a designated depot location x_0 , in which case there is no restriction on the timing of decisions.
- P2. A decision consists exclusively of choosing to visit either one of the demands in the system or to visit (or remain at) the fixed depot location x_0 .
- P3. Let $W(x)$ be the waiting time conditioned on a random arrival's location being x ; that is $W(x) \equiv E[W_i | X_i = x]$. Then the policy provides *normalized waiting times* that are bounded and smooth; specifically, the function

$$(3) \quad \psi(x) \equiv \frac{W(x)}{W}$$

satisfies $0 < \psi \leq \psi(x) \leq \bar{\psi} < \infty$ and is K -Lipschitz for all $0 < \rho < 1$ when $f(x)$ is similarly bounded and K -Lipschitz.

Several characteristics of the class \mathcal{M} are worth examining. First, we need a depot location for two reasons: to ensure that there is always a decision available at each completion epoch, and to ensure that there is some location where the vehicle can remain idle if the system is empty. The fact that a vehicle can be idle only at the depot x_0 turns out to be an important property for our subsequent analysis. Second, P3 requires that the policy provide local waiting times $W(x)$ that do not vary

abruptly with changes in location x and are bounded relative to the system-wide expected waiting time W . Again, this is needed primarily for technical reasons, though it is arguably a natural characteristic of most practical policies. Indeed, by definition (see below), $W(x) = W$ for any spatially unbiased policy and thus P3 is trivially satisfied for all such policies; however, in the spatially biased case, we must restrict \mathcal{M} to only those policies which satisfy P3. We remark that in the uniform demand case discussed in [6] and [7], we allowed a more general class of policies in which vehicles can wait at any location x and change destinations at any time and smoothness of $\psi(x)$ was not required. Finally, we note that since the velocity v is finite and the location distribution is continuous, service completions are sequential (i.e. service completions do not occur simultaneously).

Letting T_μ denote the system time of a particular policy $\mu \in \mathcal{M}$, then the DTRP is the problem of finding a policy μ^* , if one exists, such that

$$T_{\mu^*} = \inf \{ T_\mu \mid \mu \in \mathcal{M} \}.$$

We let T^* denote the infimum on the right-hand side above. A policy μ for which T_μ/T^* is bounded has a *constant factor guarantee*. If $\lim_{\rho \rightarrow 1} (T_\mu/T^*)$ is bounded, then the policy μ has a constant factor guarantee in heavy traffic.

2.2. *Spatially biased and unbiased policies.* We shall also need two definitions mentioned informally above and in Section 1. In these definitions, X_i is the location of a randomly chosen demand and W_i is its waiting time.

Definition 1. A policy $\mu \in \mathcal{M}$ is *spatially unbiased* if

$$E[W_i \mid X_i \in \mathcal{S}] = W \quad \forall \mathcal{S} \subseteq \mathcal{A}.$$

Definition 2. A policy $\mu \in \mathcal{M}$ is *spatially biased* if there exists sets $\mathcal{S}_1, \mathcal{S}_2 \subseteq \mathcal{A}$ such that

$$E[W_i \mid X_i \in \mathcal{S}_1] > E[W_i \mid X_i \in \mathcal{S}_2].$$

Recalling that $W(x) \equiv E[W_i \mid X_i = x]$, observe that for spatially unbiased policies $W = W(x)$ for all x as mentioned above. For a spatially biased policy W and $W(x)$ are related only by

$$W = \int_{\mathcal{A}} W(x)f(x) dx.$$

2.3. *Spatial queues on subsets of \mathbb{R}^2 .* We associate a queue with every $\mathcal{S} \subseteq \mathbb{R}^2$, referred to as a *spatial queue*, by considering \mathcal{S} to be a ‘black box’ that has arrivals (demands arriving to \mathcal{S}) and departures (service completions within \mathcal{S}). Let $p(\mathcal{S}) = \int_{\mathcal{S}} f(x) dx$ denote the probability that an arrival falls in the set \mathcal{S} , $\lambda(\mathcal{S}) = \lambda p(\mathcal{S})$ denote the arrival rate to \mathcal{S} , and $N(\mathcal{S})$ denote the time average number of demands in the queue \mathcal{S} . Note that since all demands are located in \mathcal{A} , $N(\mathcal{S}) = N(\mathcal{S} \cap \mathcal{A})$ for all sets \mathcal{S} . In particular, the total time average number in

queue, N , is given by $N = N(\mathcal{A})$. The waiting time in \mathcal{S} is denoted $W(\mathcal{S})$ and given by

$$W(\mathcal{S}) = E[W_i | X_i \in \mathcal{S}] = \int_{\mathcal{S}} W(x)f(x | X_i \in \mathcal{S}) dx.$$

We define

$$(4) \quad \phi(x) \equiv f(x)\psi(x).$$

The function $\phi(x)$ is interpreted as the time-average density of demands in the system in the sense that

$$(5) \quad N(\mathcal{S}) = N \int_{\mathcal{S}} \phi(x) dx,$$

for any $\mathcal{S} \subset \mathbb{R}^2$. To see this we use the definition of $\phi(x)$ and Little’s theorem, and note that

$$\begin{aligned} N(\mathcal{S}) &= \lambda(\mathcal{S})W(\mathcal{S}) \\ &= \left[\lambda \int_{\mathcal{S}} f(x) dx \right] \left[\int_{\mathcal{S}} W(x)f(x | X_i \in \mathcal{S}) dx \right] \\ &= \lambda W \int_{\mathcal{S}} \frac{W(x)}{W} f(x) dx \\ &= N \int_{\mathcal{S}} \phi(x) dx. \end{aligned}$$

3. Heavy traffic lower bounds

In this section, we derive our main lower bounds on T^* . We begin by proving some lemmas related to spatial queues in the system that culminate in a generic lower bound on T^* . We then specialize this bound to the spatially biased and unbiased cases to arrive at our main lower bounds.

3.1. Preliminary lemmas and bounds

3.1.1. *Preliminary lemmas.* Our first lemma shows that $\phi(x)$ has the interpretation of a density function and that it is both bounded and smooth.

Lemma 1. For any stable policy satisfying property P3, the function $\phi(x)$ satisfies

$$(6) \quad \int_{\mathcal{A}} \phi(x) dx = 1$$

$$(7) \quad \phi(x) \geq 0 \quad \forall x \in \mathcal{A},$$

and $\phi(x)$ is bounded, $0 < \underline{\phi} \leq \phi(x) \leq \bar{\phi} < \infty$, and K -Lipschitz.

Proof. The first part, (6)–(7), follows directly from the definition of $\phi(x)$ and the fact that $N(\mathcal{A}) = N$ and $N(\cdot)$ is always positive. Boundedness of $\phi(x)$ is due to the

boundedness of $f(x)$ and $\psi(x)$. The K -Lipschitz condition follows by noting that

$$|f(x)\psi(x) - f(y)\psi(y)| = |f(x)(\psi(x) - \psi(y)) - \psi(y)(f(y) - f(x))| \\ \leq \bar{f}K \|x - y\| + \bar{\psi}K \|x - y\| = K' \|x - y\|.$$

In our proof of the lower bounds we will need to relate the expected number of demands left behind by a random departure from a small area \mathcal{S} to the time-average number in queue $N(\mathcal{S})$. Our next lemma gives us this relationship.

Let the random variables $Y(\mathcal{S})$ denote an interarrival time of the queue \mathcal{S} . Note that arrivals to \mathcal{S} represent a thinning of the renewal process to the entire region and thus $Y(\mathcal{S})$ is a geometric sum of interarrival times. Therefore, the Laplace transform $A_1^*(s)$ of the resulting renewal process is given by

$$(8) \quad A_1^*(s) = \frac{A^*(s)p(\mathcal{S})}{1 - A^*(s)(1 - p(\mathcal{S}))},$$

where $p(\mathcal{S}) = \int_{\mathcal{S}} f(x) dx$ is the probability that an arrival falls in the set \mathcal{S} . In particular the coefficient of variation $c_{A_1}^2 = \text{Var}[A_1]/(E[A_1])^2$ is

$$(9) \quad c_{A_1}^2 = 1 + p(\mathcal{S})(c_A^2 - 1),$$

where c_A^2 is the coefficient of variation of the interarrival time of the original renewal process.

Let $n^+(\mathcal{S})$, a random variable, denote the number of demands left behind by a random departure from \mathcal{S} and $n(\mathcal{S})$ denote the number of demands in queue at a random epoch. Since service is sequential, this is a well-defined random variable. Define $N^+(\mathcal{S}) \equiv E[n^+(\mathcal{S})]$ and recall $N(\mathcal{S}) \equiv E[n(\mathcal{S})]$. Let $W(\mathcal{S})$ denote the expected waiting time in this queue. We are primarily interested in small balls and thus we define $\mathcal{C}(x, z)$ to be the set of points within a distance of z from a given location x (i.e. $\mathcal{C}(x, z) = \{y \mid \|y - x\| \leq z\}$).

If the arrival process is Poisson, then from the PASTA property (see [20]), $n^+(\mathcal{S}) =_d n(\mathcal{S})$ and thus $N^+(\mathcal{S}) = N(\mathcal{S})$. In the next lemma we address the relation of $N^+(\mathcal{S})$ to $N(\mathcal{S})$ in heavy traffic for $\mathcal{S} = \mathcal{C}(x, z)$, general renewal processes and small z . Intuitively, for small z , the probability $p(\mathcal{S})$ tends to zero and thus the coefficient of variation $c_{A_1}^2$ tends to 1, i.e. the resulting renewal process approaches a Poisson process for which we can apply PASTA. We formalize this as follows.

Lemma 2. For any stable policy satisfying P3, as $\rho \rightarrow 1$,

$$N^+(\mathcal{C}(x, z)) = N\phi(x)\pi z^2 + f(x)\frac{c_A^2 - 1}{2}\pi z^2 + No(z^2).$$

Proof. In general, the expected number of customers in the system left by a departing customer depends on the service discipline when the arrival process is not Poisson. Bertsimas and Nakazato [4] prove that in general queueing systems with renewal inputs, the expected number of customers in the system left by a departing

customer satisfies in heavy traffic:

$$N^+ = N + \frac{c_a^2 - 1}{2} + o(1),$$

where the term $o(1)$ goes to zero exponentially fast as $\rho \rightarrow 1$, and c_a^2 is the squared coefficients of variation in the input process, i.e. in heavy traffic N^+ does not depend on the service discipline. Applying this relation to the queue $\mathcal{C}(x, z)$ we obtain that

$$\begin{aligned} (10) \quad N^+(\mathcal{C}(x, z)) &= N(\mathcal{C}(x, z)) + \frac{c_{A_1}^2 - 1}{2} + o(1) \\ &= N(\mathcal{C}(x, z)) + p(\mathcal{C}(x, z)) \frac{c_A^2 - 1}{2} + o(1), \end{aligned}$$

where we have used (9).

By assumption, $f(x)$ is K -Lipschitz and thus $-K \|y - x\| + f(x) \leq f(y) \leq K \|y - x\| + f(x)$, which leads to

$$f(x)\pi z^2 - K \int_0^z r 2\pi r \, dr \leq p(\mathcal{C}(x, z)) = \int_{\mathcal{C}(x, z)} f(y) \, dy \leq f(x)\pi z^2 + K \int_0^z r 2\pi r \, dr,$$

and thus

$$f(x)\pi z^2 - \frac{2\pi}{3} K z^3 \leq p(\mathcal{C}(x, z)) \leq f(x)\pi z^2 + \frac{2\pi}{3} K z^3,$$

i.e.

$$(11) \quad p(\mathcal{C}(x, z)) = f(x)\pi z^2 + o(z^2).$$

Similarly, since $N(\mathcal{C}(x, z)) = N \int_{\mathcal{C}(x, z)} \phi(y) \, dy$, and by Lemma 1, $\phi(y)$ is K -Lipschitz,

$$(12) \quad N(\mathcal{C}(x, z)) = N\phi(x)\pi z^2 + No(z^2).$$

Substituting (11), (12) into (10) proves the lemma.

3.1.2. Preliminary lower bounds. We now combine the above results to derive an important lemma relating the expected nearest-neighbor distance at a completion epoch to N , the average number in queue. Let d_i denote the distance traveled from demand i to the next demand served after i ; that is, the distance the serving vehicle travels after *departing* from i . Let Z_i^* denote the distance from the server to either the depot or the closest unserved demand (whichever is smaller) at the completion epoch of demand i (the ‘nearest neighbor’ distance). Then we have the following result.

Lemma 3. For any stable policy satisfying P1 and P2, $E[Z_i^*] \leq E[d_i]$ and therefore

$$E[Z^*] \equiv \lim_{i \rightarrow \infty} E[Z_i^*] \leq \lim_{i \rightarrow \infty} E[d_i] \equiv \bar{d}.$$

Proof. Properties P1 and P2 imply that Z_i^* is the decision that minimizes d_i for all i and thus the above follows.

Finding a bound on $E[Z^*]$ will thus provide a bound on $E[d]$. Such a bound is provided by the following key lemma.

Lemma 4. For any stable policy satisfying P3,

$$\lim_{N \rightarrow \infty} \sqrt{N} E[Z^*] \geq \frac{2}{3\sqrt{\pi}} \int_{\mathcal{A}} \phi^{-\frac{1}{2}}(x) f(x) dx.$$

Proof. Consider a randomly tagged demand arriving at location X and condition on the event $\{X = x\}$. Define $c(x) \equiv N\pi\phi(x)$ and let $\mathcal{D}(N) \equiv \{x \mid \|x - x_0\| \leq c^{-\frac{1}{2}}(x)\}$ denote the set of points x within a radius $c^{-\frac{1}{2}}(x)$ of the location x_0 . There are two cases to consider:

Case 1: $x \notin \mathcal{D}(N)$. Recall that $n^+(\mathcal{C}(x, z))$ denotes the number of demands in the set $\mathcal{C}(x, z)$ at the completion epoch of our tagged demand, and observe that as long as the depot x_0 is outside the ball $\mathcal{C}(x, z)$,

$$(13) \quad P(Z^* \leq z \mid X = x) = P(n^+(\mathcal{C}(x, z)) > 0) \leq N^+(\mathcal{C}(x, z)),$$

where the last inequality is due to the fact that $n^+(\mathcal{C}(x, z))$ is a non-negative, integer-valued random variable. Considering the service completion of our tagged demand as a departure from the queue $\mathcal{C}(x, z)$, we therefore have by Lemma 2 and the boundedness of $f(x)$ that

$$N^+(\mathcal{C}(x, z)) = N\phi(x)\pi z^2 + O(1)z^2 + No(z^2).$$

Substituting into the bound (13) implies

$$P(Z^* > z \mid X = x) \geq 1 - N\phi(x)\pi z^2 - O(1)z^2 - No(z^2).$$

Recalling that $c(x) \equiv N\pi\phi(x)$, and the depot is outside the ball $\mathcal{C}(x, z)$ for all $z \leq c^{-\frac{1}{2}}(x)$ when $x \notin \mathcal{D}(N)$ we have

$$\begin{aligned} E[Z^* \mid X = x] &= \int_0^\infty P(Z^* > z \mid X = x) dz \\ &\geq \int_0^{c^{-\frac{1}{2}}(x)} \max\{0, 1 - N\pi z^2 \phi(x) - O(1)z^2 - No(z^2)\} dz \\ &\geq \int_0^{c^{-\frac{1}{2}}(x)} (1 - c(x)z^2) dz - \int_0^{c^{-\frac{1}{2}}(x)} (O(1)z^2 + No(z^2)) dz \\ &= \frac{2}{3\sqrt{\pi N}} \phi^{-\frac{1}{2}}(x) - o(N^{-\frac{1}{2}}). \end{aligned}$$

Case 2: $x \in \mathcal{D}(N)$. In this case $n^+(\mathcal{C}(x, z)) = 0$ does not imply $Z^* > z$ for all $z \leq c^{-\frac{1}{2}}(x)$. Thus we use only the trivial bound

$$P(Z^* > z \mid X = x) \geq 0.$$

We now combine these two cases by unconditioning. First note that by Lemma 1, $\phi(x) \geq \phi \forall N$, which implies $\int_{\mathcal{D}(N)} dx \leq O(1/N)$. Therefore,

$$\begin{aligned} E[Z^*] &\geq \frac{2}{3\sqrt{\pi N}} \int_{\mathcal{A}-\mathcal{D}(N)} \phi^{-\frac{1}{2}}(x)f(x) dx - o(N^{-\frac{1}{2}}) \\ &\geq \frac{2}{3\sqrt{\pi N}} \left[\int_{\mathcal{A}} \phi^{-\frac{1}{2}}(x)f(x) dx - \phi^{-\frac{1}{2}}\bar{f} \int_{\mathcal{D}(N)} dx \right] - o(N^{-\frac{1}{2}}) \\ &= \frac{2}{3\sqrt{\pi N}} \int_{\mathcal{A}} \phi^{-\frac{1}{2}}(x)f(x) dx - o(N^{-\frac{1}{2}}). \end{aligned}$$

Multiplying both sides above by \sqrt{N} and taking the limit as $N \rightarrow \infty$ proves the lemma.

Remark. The proof above shows the contribution to the lower bound from the points $\mathcal{D}(N)$ that are close to the depot is asymptotically insignificant. Thus, though we require a depot to ensure that decisions are well defined at each completion epoch, its presence has no effect on the heavy traffic bounds.

Lemma 4 can be used to prove the following intermediate bound on the optimal system time T^* .

Lemma 5. There exists a single constant γ such that for all stable policies satisfying P1–P3,

$$\lim_{\rho \rightarrow 1} T(1 - \rho)^2 \geq \gamma^2 \frac{\lambda \left[\int_{\mathcal{A}} \phi^{-\frac{1}{2}}(x)f(x) dx \right]^2}{v^2 m^2},$$

and $\gamma \geq 2/3\sqrt{\pi}$.

Proof. Consider the following necessary condition for stability

$$\bar{s} + \frac{\bar{d}}{v} \leq \frac{m}{\lambda}.$$

Using the fact that $E[Z^*] \leq \bar{d}$ for policies that satisfy P1 and P2 by Lemma 3, multiplying the second term on the left-hand side above by \sqrt{N}/\sqrt{N} and rearranging implies

$$\sqrt{N}(1 - \rho) \geq \frac{\lambda\sqrt{N}E[Z^*]}{mv}.$$

Note that since N is at least as large as the mean number in queue in the corresponding $G/G/m$ queue (i.e. the queue with $v = \infty$), as $\rho \rightarrow 1$, we must have $N \rightarrow \infty$. Therefore taking the limit as $\rho \rightarrow 1$ (and consequently $N \rightarrow \infty$) on both sides

above and applying Lemma 4 we obtain

$$\lim_{\rho \rightarrow 1} \sqrt{N} (1 - \rho) \geq \gamma \frac{\lambda \int_{\mathcal{A}} \phi^{-\frac{1}{2}}(x) f(x) dx}{mv},$$

where $\gamma \geq 2/3\sqrt{\pi}$. Squaring both sides and using $T \geq W = N/\lambda$ completes the proof.

3.2. *A spatially unbiased lower bound.* As mentioned, Lemma 5 is only an intermediate bound since the functions $\phi(x)$ remains unspecified. Determining $\phi(x)$ for the unbiased case gives us the first of our main heavy traffic theorems.

Theorem 1. Within the class of spatially unbiased policies satisfying P1–P3,

$$\lim_{\rho \rightarrow 1} T^*(1 - \rho)^2 \geq \gamma^2 \frac{\lambda \left[\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx \right]^2}{m^2 v^2}$$

where $\gamma \geq 2/3\sqrt{\pi}$.

Proof. For a spatially unbiased policy, $W(x) = W, \forall x \in \mathcal{A}$. Thus, $\psi(x) = W(x)/W = 1$ and $\phi(x) = \psi(x)f(x) = f(x)$. Substituting $\phi(x) = f(x)$ into Lemma 5 we obtain Theorem 1.

This theorem differs from the heavy traffic bound in [6] and [7] in several ways. In one sense it is weaker because it is an asymptotic bound while our earlier bounds are valid for all values of $\rho < 1$. However, it has several important advantages. First, it applies to a general density $f(x)$ and general arrival process rather than just to the Poisson, uniform case. Second, it improves on the constant value γ by a factor of $\sqrt{2}$ and thus increases over previous lower bound by a factor of 2.

3.3. *A spatially biased lower bound.* Theorem 1 gives an asymptotic bound for the case where unbiased service is a constraint, perhaps imposed as a matter of policy. What is the system time behavior when this constraint is relaxed? The answer, in part, is provided by our second main theorem.

Theorem 2. Within the class of spatially biased policies satisfying P1–P3,

$$\lim_{\rho \rightarrow 1} T^*(1 - \rho)^2 \geq \gamma^2 \frac{\lambda \left[\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx \right]^3}{m^2 v^2}$$

where $\gamma \geq 2/3\sqrt{\pi}$.

Proof. Since no assumption of unbiased service is made, we cannot determine $\phi(x)$ as easily. However, consider the following minimization problem for the

integral term in Lemma 5:

$$\begin{aligned}
 (14) \quad z^* &= \min \int_{\mathcal{A}} \phi^{-\frac{1}{2}}(x) f(x) dx \\
 &\text{subject to } \int_{\mathcal{A}} \phi(x) dx = 1 \\
 &\phi(x) \geq 0.
 \end{aligned}$$

Using the value z^* as a lower bound on the integral term in Lemma 5 will give us Theorem 1.

Note that the objective function is convex in $\phi(x)$ and the constraints are linear; thus, (14) is a convex program. Relaxing the equality constraint above with a multiplier, we obtain the following Lagrangian dual:

$$\begin{aligned}
 (15) \quad z^*(\mu) &= \min_{\phi(x) \geq 0} \int_{\mathcal{A}} \phi^{-\frac{1}{2}}(x) f(x) dx + \mu \left[\int_{\mathcal{A}} \phi(x) dx - 1 \right] \\
 &= \int_{\mathcal{A}} \min_{\phi(x) \geq 0} [\phi^{-\frac{1}{2}}(x) f(x) + \mu \phi(x)] dx - \mu.
 \end{aligned}$$

By differentiating the integrand above and setting it equal to zero, we see that a pair $(\phi^*(x), \mu^*)$ for which

$$(16) \quad -\frac{1}{2}[\phi^*(x)]^{-\frac{3}{2}} f(x) + \mu^* = 0 \quad \forall x \in \mathcal{A}$$

$$(17) \quad \int_{\mathcal{A}} \phi^*(x) dx = 1$$

$$(18) \quad \phi^*(x) \geq 0$$

will satisfy the Kuhn–Tucker necessary conditions for optimality. One can verify by substitution that

$$(19) \quad \phi^*(x) = \left[\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx \right]^{-1} f^{\frac{2}{3}}(x)$$

$$(20) \quad \mu^* = \frac{1}{2} \left[\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx \right]^{\frac{3}{2}}$$

is such a pair. The fact that (14) is a convex program implies that these conditions are also sufficient to assure global optimality. Substituting the value $\phi^*(x)$ above into Lemma 5 gives us the theorem.

Remark. Note that even though we did not require it, the resulting $\phi(x)$ above is bounded since $f(x)$ is bounded and is K -Lipschitz since

$$\begin{aligned}
 |f^{\frac{2}{3}}(x) - f^{\frac{2}{3}}(y)| &= f^{-\frac{1}{3}}(x) f(x) - f^{-\frac{1}{3}}(y) f(y) \\
 &\leq |f^{-\frac{1}{3}}(y) f(x) - f^{-\frac{1}{3}}(y) f(y)| \leq f^{-\frac{1}{3}} |f(x) - f(y)| \leq K' \|x - y\|,
 \end{aligned}$$

where we have assumed without loss of generality that $f(x) \geq f(y)$.

4. Heavy traffic policies

We next examine two policies that have provably good performance with respect to the lower bounds of Theorems 1 and 2. The policies are modifications of policies introduced in [6] and [7].

4.1. *A provably good spatially unbiased policy.* The spatially unbiased policy we consider is defined as follows.

The unbiased (U) TSP policy. Let k be a fixed positive integer. From a central point in the interior of \mathcal{A} , subdivide the service region into k wedges $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_k$ such that $\int_{\mathcal{A}_i} f(x) dx = 1/k, i = 1, 2, \dots, k$. (One could do this by ‘sweeping’ the region from the depot using an arbitrary starting ray until $\int_{\mathcal{A}_1} f(x) dx = 1/k$, continuing the sweep until $\int_{\mathcal{A}_2} f(x) dx = 1/k$, etc.) Within each subregion, form sets of size n/k (n is a parameter to be determined.) As sets are formed, deposit them in a queue and service them FCFS with the first available vehicle by solving a TSP on the set and following it in an arbitrary direction. Optimize over n .

The following theorem shows that this unbiased policy is guaranteed to be within about 80% of the optimal policy in heavy traffic.

Theorem 3. Let T_U^* be the optimal system time over the class of spatially unbiased policies satisfying P1–P3. Then

$$\frac{T_U}{T_U^*} \leq \frac{\beta^2}{2\gamma^2} \approx 1.8 \quad \text{as } \rho \rightarrow 1.$$

where $\beta \approx 0.72$ is the TSP constant in the Euclidean plane (see [3] and [11]).

Proof. We first obtain some moments for the random variable τ , the time to service a set. Let L_i denote the length of the optimal TSP on a set in region i . Note that

$$E[\tau] = \frac{n}{k} \bar{s} + \frac{1}{v} \sum_{i=1}^k \frac{1}{k} E[L_i].$$

Observe that $kf(x)$ is the conditional density in any given subregion. From the asymptotic TSP results of [3] and [19], we can therefore assert that almost surely

$$\lim_{n \rightarrow \infty} \frac{L_i}{\sqrt{n}} = \beta \int_{\mathcal{A}_i} f^{\frac{1}{2}}(x) dx$$

and that $E[L_i]/\sqrt{n}$ converges to this value as well. Thus,

$$\begin{aligned} \frac{E[\tau]}{(n/k)} &= \bar{s} + \frac{1}{v\sqrt{n}} \sum_{i=1}^k \frac{E[L_i]}{\sqrt{n}} \\ (21) \quad &\sim \bar{s} + \frac{1}{v\sqrt{n}} \sum_{i=1}^k \beta \int_{\mathcal{A}_i} f^{\frac{1}{2}}(x) dx \\ &= \bar{s} + \frac{\beta}{v\sqrt{n}} \int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx. \end{aligned}$$

To determine σ_τ^2 , consider the random variable L which is an equiprobable selection from the set of random variables $\{L_1, \dots, L_k\}$. That is, L is the random variable such that

$$\sigma_\tau^2 = \frac{n}{k} \sigma_s^2 + \frac{1}{v^2} \text{Var} [L] = \frac{n}{k} \sigma_s^2 + \frac{1}{k} \sum_{i=1}^k \text{Var} L_i.$$

Since $\text{Var} L_i = O(1)(1/k)$ (see Karp and Steele [13]) we obtain

$$(22) \quad \frac{\sigma_\tau^2}{(n/k)} = \sigma_s^2 + O\left(\frac{1}{n}\right).$$

We will use (21) and (22) shortly.

Note that each region independently generates its own arrival stream of sets and thus the input to the resulting queue of sets is the superposition of k renewal processes, one from each region. (A queue is denoted $\Sigma GI/G/m$ if its input process is the superposition of k independent renewal processes (not necessarily identical).) We analyze this resulting queue using the following theorem of Iglehart and Whitt.

Theorem 4 (Iglehart and Whitt [10]). Consider an m -server queue fed by the superposition of k renewal processes. Let $1/\lambda_i$ and $\sigma_{a_i}^2$ denote, respectively, the mean and variance of the interarrival time of the i th renewal process, $i = 1, 2, \dots, k$. Let $1/\mu_j$ and $\sigma_{b_j}^2$ denote the mean and variance, respectively, of the service times at server $j = 1, 2, \dots, m$. Define $\lambda \equiv \sum_{i=1}^k \lambda_i \mu \equiv \sum_{j=1}^m \mu_j$ and $\rho \equiv \lambda/\mu$. Then as $\rho \rightarrow 1$ the mean waiting time in queue, W , satisfies

$$(23) \quad W \sim \frac{\sum_{i=1}^k \lambda_i^3 \sigma_{a_i}^2 + \sum_{j=1}^m \mu_j^3 \sigma_{b_j}^2}{2\mu^2(1-\rho)}.$$

Let $\hat{\lambda}_i = \lambda/n$ denote the arrival rate of sets to region i and $\hat{\lambda} = \sum_{i=1}^k \hat{\lambda}_i = k\lambda/n$ denote the overall arrival rate of sets. Since the interarrival time in each subregion is a geometric sum of interarrival times in the entire region, one can easily show that the variance of the interarrival time of sets from subregion i , $\sigma_{a_i}^2$, is given by

$$\sigma_{a_i}^2 = n \left(\sigma_a^2 + \frac{k-1}{\lambda^2} \right),$$

where σ_a^2 is the variance of the interarrival times of demands to the entire region \mathcal{A} . This implies that

$$\sum_{i=1}^k \hat{\lambda}_i^3 \sigma_{a_i}^2 = \hat{\lambda}^2 \left(\frac{\lambda \sigma_a^2}{k} + \frac{1-(1/k)}{\lambda} \right).$$

We shall use the fact that for large values of k , the right-hand side above is approximately $(1/\lambda)\hat{\lambda}^2$, and in heavy traffic, $\hat{\lambda} \approx m/E[\tau]$. Using these facts and applying (21) and (22), we can therefore establish the following limit for W_{set} , the

time a set waits in queue:

$$W_{\text{set}} \sim \frac{\lambda \left(\frac{1}{\lambda^2} + \frac{1}{m^2} \frac{\sigma_\tau^2}{(n/k)} \right)}{2 \left(1 - \frac{k\lambda}{mn} E[\tau] \right)} \sim \frac{\lambda \left[\frac{1}{\lambda^2} + \frac{1}{m^2} \left(\sigma_s^2 + O\left(\frac{1}{n}\right) \right) \right]}{2 \left(1 - \rho - \frac{\lambda\beta}{mv\sqrt{n}} \int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx \right)}$$

Note that the stability condition for this queue implies that

$$\rho + \frac{\lambda\beta}{mv\sqrt{n}} \int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx < 1,$$

which implies

$$n > \frac{\lambda^2 \beta^2 \left(\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx \right)^2}{m^2 v^2 (1 - \rho)^2},$$

so $n \rightarrow \infty$ as $\rho \rightarrow 1$ and thus using TSP asymptotics is valid in heavy traffic.

The waiting time W_{set} is not itself the wait for service of an individual demand; it is the wait in queue for a set. The time of arrival of a set is actually the time of arrival of the last demand in that set. Therefore, we must add the time a demand waits for its set to form, denoted W^- , and also the time it takes to complete service of the demand once its set enters service, denoted W^+ . By conditioning on the position of a randomly chosen demand in its tour, one can easily show that

$$W^- \leq \frac{1}{2} \left(\frac{n}{k} \right) \frac{k}{\lambda} = \frac{n}{2\lambda},$$

and

$$W^+ \leq \frac{1}{2} \left(\frac{n}{k} \right) \bar{s} + O(\sqrt{n})$$

where the $O(\sqrt{n})$ term is due to the TSP travel cost to service the sets of size n/k .

Adding W^- , W^+ and W_{set} we obtain the following bound on T_U :

$$T_U \leq \frac{n \left(1 + \frac{m\rho}{k} \right)}{2\lambda} + \frac{\lambda \left[\frac{1}{\lambda^2} + \frac{1}{m^2} \left(\sigma_s^2 + O\left(\frac{1}{n}\right) \right) \right]}{2 \left(1 - \rho - \frac{\lambda\beta}{mv\sqrt{n}} \int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx \right)} + O(\sqrt{n}).$$

Making a change of variable to

$$y = \frac{\lambda\beta \int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx}{mv(1 - \rho)\sqrt{n}},$$

the bound can be written

$$T_U \leq \frac{\lambda\beta^2 \left(\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx \right)^2 \left(1 + \frac{m\rho}{k} \right)}{2m^2 v^2 (1 - \rho)^2 y^2} + \frac{\lambda \left[\frac{1}{\lambda^2} + \frac{1}{m^2} \left(\sigma_s^2 + O(y^2(1 - \rho)^2) \right) \right]}{2(1 - \rho)(1 - y)} + O\left(\frac{1}{y(1 - \rho)}\right).$$

An approximate optimal value for y is

$$y^* \approx 1 - \frac{mv}{\beta} \sqrt{\frac{\left[\frac{1}{\lambda^2} + \frac{\sigma_s^2}{m^2}\right](1-\rho)}{2\left(\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx\right)^2\left(1 + \frac{m\rho}{k}\right)}}.$$

Substituting this value into the above bound we find that as $\rho \rightarrow 1$,

$$T_U \sim \frac{\lambda\beta^2\left(\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx\right)^2\left(1 + \frac{m}{k}\right)}{2m^2v^2(1-\rho)^2},$$

where the second-order term is $O((1-\rho)^{-\frac{3}{2}})$. The theorem then follows by comparing the above leading behavior to the bound in Theorem 1 and choosing k arbitrarily large.

Remark. Note the k radial cuts play an important role in this policy. They reduce the third component of the waiting time, W^+ , the wait for the vehicle to reach a demand once its set enters service, by a factor of $1/k$ over our previous policies in [6] which used no radial cuts. Since this component of the waiting time, along with W^+ , is a leading order term, the reduction is significant. Indeed, this minor modification results in a system time one half as large as the original TSP policy presented in [6].

Remark. Though the provable guarantee on this policy is 1.8, we conjecture that it is in fact an asymptotically optimal policy in heavy traffic. We base this conjecture on heuristic arguments. In particular, we have been able to show that if it is optimal to always clear a region (such as a radial region) before moving on to service a new region, then $\gamma = \beta/\sqrt{2}$ in Theorem 1. A complete proof of this conjecture, however, remains an interesting challenge.

4.2. *A provably good spatially biased policy for piecewise uniform demand.* We next propose a policy that achieves a performance guarantee of $\beta^2/2\gamma^2 \approx 1.8$ with respect to the spatially biased bound lower bound when f is a *piecewise uniform density*, i.e. there exists a partition of \mathcal{A} into J subsets $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_J$ with bounded perimeter P such that $f(x) = \mu_j \forall x \in \mathcal{A}_j, j = 1, 2, \dots, J$. For such a density

$$\int_{\mathcal{A}} f^{\frac{3}{2}}(x) dx = \sum_{j=1}^J \mu_j^{\frac{3}{2}} A_j.$$

Note that this density is not K -Lipschitz; however, one can modify it by linearly smoothing the function throughout a strip of width $\varepsilon > 0$ about the perimeter between the subsets $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_J$. Call this smoothed function $f_\varepsilon(x)$. The resulting $f_\varepsilon(x)$ is K -Lipschitz with K of order $(1/\varepsilon) \max_{i,j} |\mu_i - \mu_j|$. Further, since the

area over which $f(x) \neq f_\epsilon(x)$ is $O(P\epsilon)$ and $|f^{\frac{2}{3}}(x) - f_\epsilon^{\frac{2}{3}}(x)| = O(\max_{i,j} |\mu_i^{\frac{2}{3}} - \mu_j^{\frac{2}{3}}|)\epsilon$,

$$\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx = \int_{\mathcal{A}} f_\epsilon^{\frac{2}{3}}(x) dx + O\left(P\left(\max_{i,j} |\mu_i^{\frac{2}{3}} - \mu_j^{\frac{2}{3}}|\right)\epsilon\right),$$

and thus $|\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx - \int_{\mathcal{A}} f_\epsilon^{\frac{2}{3}}(x) dx| = O(\epsilon)$. In what follows we analyze the problem and state our results for the piecewise uniform case directly; however, one should bear in mind that we can always approximate such a $f(x)$ arbitrarily closely by a K -Lipschitz function $f_\epsilon(x)$ in order to satisfy the technical requirements of Theorem 1.

Though such piecewise constant densities are not perfectly general, one could approximate a continuous density by a piecewise constant density. Moreover, in practice a piecewise uniform density is probably an adequate model.

The policy itself is defined as follows.

The biased (B) TSP policy. Let $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_J$ be a partition of \mathcal{A} such that $f(x) = \mu_j \forall x \in \mathcal{A}_j, j = 1, 2, \dots, J$. Let A_j denote the area of \mathcal{A}_j . For a given positive integer k , partition each subset \mathcal{A}_j further into $k_j = \mu_j^{\frac{2}{3}} A_j k$ regions of area $A_j/k_j = (\mu_j^{\frac{2}{3}} k)^{-1}$ (k is a scale factor that will be chosen arbitrarily large; hence, we assume an integer k_j can be found such that k_j/k is sufficiently close to $\mu_j^{\frac{2}{3}} A_j$). Within each of these subregions, form demands into sets of size n/k as they arrive. As sets are formed, deposit them in a queue and service them FCFS with the first available vehicle as follows: (1) solve a TSP on the set; (2) connect the tour to the depot through an arbitrary point in the tour; and (3) follow the resulting tour in an arbitrary direction servicing demands as they are encountered. Optimize over n .

Let the system of this policy be denoted T_B . We shall prove the following theorem.

Theorem 5. If f is a piecewise uniform density and T_B^* is the optimal system time over the class of biased policies satisfying P1–P3, then

$$\frac{T_B}{T_B^*} \leq \frac{\beta^2}{2\gamma^2} \approx 1.8 \quad \text{as } \rho \rightarrow 1.$$

Proof. We again begin by obtaining the first two moments of the random variable τ , the time to service a randomly chosen set of demands. A set formed in \mathcal{A}_j will be called a *type j set*. Let $p_j \equiv \mu_j A_j$ denote the probability that a randomly selected set is a type j set. (Note that since the set size is n/k in all subregions, the probability that a randomly selected demand is contained in a type j set is the same as the probability that a randomly selected set is of type j .) Let the random variable L_j denote the length of a tour on a type j set. Then

$$E[\tau] = (n/k)\bar{s} + \frac{1}{v} \sum_{j=1}^J p_j E[L_j].$$

We show below that as $\rho \rightarrow 1$, $n \rightarrow \infty$; therefore

$$\frac{kE[L_j]}{(\sqrt{n})} \rightarrow \sqrt{k}\beta \sqrt{\frac{A_j}{k_j}} = \beta\mu_j^{-\frac{1}{3}}.$$

Note that the connection cost to the depot is $O(1)$ and thus its contribution to $E[L_j]/\sqrt{n}$ is negligible as $n \rightarrow \infty$. Substituting this above implies that as $n \rightarrow \infty$

$$(24) \quad \frac{E[\tau]}{(n/k)} \rightarrow \bar{s} + \frac{\beta}{v\sqrt{n}} \sum_{j=1}^J p_j \mu_j^{-\frac{1}{3}} = \bar{s} + \frac{\beta}{v\sqrt{n}} \sum_{j=1}^J \mu_j^{\frac{2}{3}} A_j.$$

To determine σ_τ^2 we let L be a random variable such that $L = L_j$ with probability p_j , $j = 1, \dots, J$. Then

$$\sigma_\tau^2 = \frac{n}{k} \sigma_s^2 + \text{Var}[L] = \frac{n}{k} \sigma_s^2 + \sum_{j=1}^J p_j \text{Var}[L_j].$$

Since $\text{Var}[L_j] = O(1)$, for large n and fixed k

$$(25) \quad \frac{\sigma_\tau^2}{n/k} = \sigma_s^2 + \frac{O(1)}{n}.$$

Defining W^- , W^+ and W_{set} as before, we have

$$(26) \quad \begin{aligned} W^- &\leq \sum_{j=1}^J p_j \frac{1}{2} \left(\frac{n/k}{(p_j \lambda)/k_j} \right) \\ &= \frac{n}{2\lambda} \sum_{j=1}^J \frac{k_j}{k} \\ &= \frac{n}{2\lambda} \sum_{j=1}^J \mu_j^{\frac{2}{3}} A_j \end{aligned}$$

and

$$(27) \quad \begin{aligned} W^+ &\leq \frac{n}{2k} \bar{s} + \frac{1}{v} \sum_{j=1}^J p_j E[L_j] \\ &= \frac{n}{2k} \bar{s} + O(\sqrt{n}). \end{aligned}$$

The queue defined by this policy is again a $\Sigma GI/G/m$ queue. Let

$$\hat{\lambda}_{ij} = \frac{k\lambda p_j}{n k_j}$$

and $\sigma_{a_{ij}}^2$, denote, respectively, the arrival rate and variance of the interarrival time of sets from the i th subregion of \mathcal{A}_j , $i = 1, \dots, k_j$. Let $\hat{\lambda} = \sum_{j=1}^J \sum_{i=1}^{k_j} \hat{\lambda}_{ij} = k\lambda/n$ denote the overall arrival rate of sets. Then, by the same reasoning as in the unbiased case we find that

$$\sigma_{a_{ij}}^2 = \frac{n}{k} \left(\frac{k_j}{p_j} \sigma_a^2 + \frac{1 - p_j/k_j}{(p_j/k_j)^2 \lambda^2} \right),$$

and therefore

$$\begin{aligned} \sum_{j=1}^J \sum_{i=1}^{k_j} \hat{\lambda}_{ij}^3 \sigma_{a_{ij}}^2 &= \sum_{j=1}^J \sum_{i=1}^{k_j} \left(\frac{k\lambda p_j}{n k_j}\right)^3 \left(\frac{n}{k}\right) \left(\frac{k_j}{p_j} \sigma_a^2 + \frac{1-p_j}{\left(\frac{p_j}{k_j}\right)^2 \lambda^2}\right) \\ &= \hat{\lambda}^2 \lambda \sum_{j=1}^J \left(p_j \left(\frac{p_j}{k_j}\right) \sigma_a^2 + \frac{1}{\lambda^2} p_j \left(1 - \frac{p_j}{k_j}\right)\right) \\ &= \hat{\lambda}^2 \left(\frac{1}{\lambda} + \frac{\lambda}{k} \left(\sigma_a^2 - \frac{1}{\lambda^2}\right) \sum_{j=1}^J \mu_j^3 A_j\right). \end{aligned}$$

Again, we use the fact that for large values of k , the right-hand side above is approximately $(1/\lambda)\lambda^2$. Substituting this approximate expression into Theorem 4 and using Equations (24) and (25) we obtain

$$W_{\text{set}} \sim \frac{\lambda \left(\frac{1}{\lambda^2} + \frac{1}{m^2} \frac{\sigma_s^2}{(n/k)}\right)}{2 \left(1 - \frac{k\lambda}{mn} E[\tau]\right)} \sim \frac{\lambda \left(\frac{1}{\lambda^2} + \frac{1}{m^2} \left(\sigma_s^2 + \frac{O(1)}{n}\right)\right)}{2 \left(1 - \rho - \frac{\lambda\beta}{mv\sqrt{n}} \sum_{j=1}^J \mu_j^3 A_j\right)}.$$

Adding the bounds (26) and (27) to the above expression we obtain that as $\rho \rightarrow 1$,

$$T_B \leq \frac{n}{2\lambda} \left(\sum_{j=1}^J \mu_j^3 A_j + \frac{m\rho}{k}\right) + \frac{\lambda \left(\frac{1}{\lambda^2} + \frac{1}{m^2} \left(\sigma_s^2 + \frac{O(1)}{n}\right)\right)}{2 \left(1 - \rho - \frac{\lambda\beta}{mv\sqrt{n}} \sum_{j=1}^J \mu_j^3 A_j\right)} + O(\sqrt{n}).$$

In terms of

$$y = \frac{\lambda\beta \sum_{j=1}^J \mu_j^3 A_j}{mv(1-\rho)\sqrt{n}}$$

we have

$$\begin{aligned} T_B &\leq \frac{\lambda\beta^2 \left[\left(\sum_{j=1}^J \mu_j^3 A_j\right)^3 + \left(\sum_{j=1}^J \mu_j^3 A_j\right)^2 \left(\frac{m\rho}{k}\right)\right]}{2m^2v^2(1-\rho)^2y^2} \\ &\quad + \frac{\lambda \left(\frac{1}{\lambda^2} + \frac{1}{m^2} \left(\sigma_s^2 + \frac{O(1)}{n}\right)\right)}{2(1-\rho)(1-y)} + O\left(\frac{1}{y(1-\rho)}\right). \end{aligned}$$

An approximate optimal value for y is

$$y^* \approx 1 - \frac{mv}{\beta} \sqrt{\frac{\left(\frac{1}{\lambda^2} + \frac{\sigma_s^2}{m^2} (1-\rho)\right)}{2 \left[\left(\sum_{j=1}^J \mu_j^3 A_j\right)^3 + \left(\sum_{j=1}^J \mu_j^3 A_j\right)^2 \frac{m\rho}{k}\right]}}.$$

Substituting this into the bound on T_B for $\rho \rightarrow 1$ we obtain

$$T_B \sim \frac{\lambda\beta^2 \left[\left(\sum_{j=1}^J \mu_j^3 A_j \right)^3 + \left(\sum_{j=1}^J \mu_j^3 A_j \right)^2 \frac{m}{k} \right]}{2m^2 v^2 (1 - \rho)^2},$$

where the second-order term is $O((1 - \rho)^{-\frac{3}{2}})$. For large k , this is arbitrarily close to

$$T_B \sim \frac{\lambda\beta^2 \left(\sum_{j=1}^J \mu_j^3 A_j \right)^3}{2m^2 v^2 (1 - \rho)^2}.$$

Comparing this to the lower bound in Theorem 2 establishes the theorem provided that the policy is in the class \mathcal{M} when $f(x)$ is replaced by a K -Lipschitz approximation $f_\epsilon(x)$. As described, the policy is not in \mathcal{M} since the waiting time $W(x)$ is discontinuous at the boundaries of the subregions; however, one can modify the policy by randomly assigning points to adjacent subregions over a strip of width ϵ about the boundary in such a way that the waiting times $W(x)$ become K -Lipschitz. Further, one can show that the resulting system time for such a modified policy satisfies

$$T_B \sim \frac{\lambda\beta^2 \left(\sum_{j=1}^J \mu_j^3 A_j + O(\epsilon) \right)^3}{2m^2 v^2 (1 - \rho)^2}.$$

Comparing this to the lower bound in Theorem 2 when $f(x)$ is replaced by $f_\epsilon(x)$ for an arbitrarily small $\epsilon > 0$ completes the proof.

4.3. *A numerical investigation of the performance of the space-filling curve (SFC) and nearest-neighbor (NN) policies.* In two policies we have examined thus far all have constant factor guarantees in heavy traffic. However, they use optimal tours, which in practice are difficult to compute. In this section, we perform simulation experiments on two alternative policies that are more computationally efficient. The first is the space-filling curve (SFC) proposed originally by Bartholdi and Platzman in [2]; the second is the simple nearest-neighbor (NN) policy. These policies show new interesting behavior and also suggest a generic approach using simulation to estimate the behavior of policies which cannot be rigorously analyzed.

The policies are defined formally as follows.

The SFC policy. Let $\mathcal{C} = \{\theta \mid 0 \leq \theta \leq 1\}$ denote the unit circle and $\mathcal{S} = \{(x, y) \mid 0 \leq x \leq 1, 0 \leq y \leq 1\}$ denote the unit square. A space-filling curve is a continuous mapping ψ from \mathcal{C} onto \mathcal{S} that preserves certain ‘nearness’ properties (cf. Platzman and Bartholdi [16] and [1] for details). The particular curve we use is defined in [1]. Without loss of generality, suppose we scale distance so that the service region \mathcal{A} is contained in \mathcal{S} and maintain the pre-images of all demands in the system (i.e. their corresponding positions in \mathcal{C}). Then the SFC policy is to service demands as they are encountered in repeated clockwise sweeps of the circle \mathcal{C} .

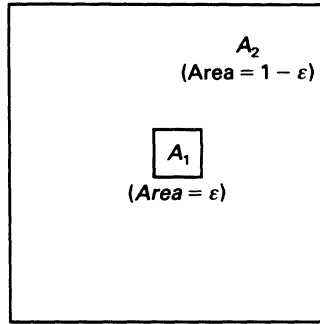


Figure 1. An extreme case general demand example

The NN policy. At each service completion epoch, the vehicle chooses to visit next the closest unserved demand.

In [6], we showed via simulation study that the system time for these two policies has the same form as the TSP policies; namely, for the uniform case,

$$T_\mu \approx \gamma_\mu^2 \frac{\lambda A}{v^2(1 - \rho)^2},$$

where $\gamma_{\text{SFC}} \approx 0.66$ and $\gamma_{\text{NN}} \approx 0.64$. In comparison, the modified TSP policy which has $\gamma_{\text{TSP}} \approx 0.51$ and the lower bound of Theorem 1 has a value $\gamma_{\text{LB}} \approx 0.38$. We next investigate the *distributional* behavior of these two policies using a similar set of simulation experiments. This analysis also suggests a generic approach to estimating the behavior of policies that cannot be rigorously analyzed.

4.3.1. *Simulation experiments.* The general demand distribution used in the simulation experiments is the one shown in Figure 1. The regions \mathcal{A}_1 and \mathcal{A}_2 have areas ε and $1 - \varepsilon$ respectively. Within each region demands are uniformly distributed. Points fall in region \mathcal{A}_1 with probability $1 - \delta$ and in region \mathcal{A}_2 with probability δ . Thus, the density is piecewise uniform with

$$(28) \quad f(x) = \begin{cases} \frac{1 - \delta}{\varepsilon} & x \in \mathcal{A}_1 \\ \frac{\delta}{1 - \varepsilon} & x \in \mathcal{A}_2. \end{cases}$$

We used identical simulation techniques (i.e. same simulation code with different $f(x)$) as in [6]. (See [6] for details.) To estimate the dependence of the system time for each policy, we set $\varepsilon = 10^{-4}$ and fixed $A = 1$, $\bar{s} = 0.1$, $\sigma_s^2 = 0$ and $\rho = 0.8$. Then, a different simulation run was performed for 11 values of δ in the range 0.05 to 0.9999. (This last value corresponds to uniform demand.) The observed average number in the system (which is proportional to the average system time) was recorded for each δ for both the SFC and NN policies.

4.3.2. *Distributional behavior of SFC and NN policies.* Before examining the results of the simulation runs, it is useful to consider the following representation of the dependence of the system time on the density $f(x)$:

$$T = \Theta \left(\frac{\lambda \Xi(\alpha)}{m^2 v^2 (1 - \rho)^2} \right)$$

where

$$\Xi(\alpha) = \left[\int_{\mathcal{A}} f^\alpha(x) dx \right]^{1/(1-\alpha)}.$$

In the unbiased case $\alpha = \frac{1}{2}$ and in the optimal biased case $\alpha = \frac{2}{3}$. For the particular density $f(x)$ given by (28) and for ϵ small,

$$\Xi(\alpha) \approx [\delta^\alpha (1 - \epsilon)^{1-\alpha}]^{1/(1-\alpha)},$$

and therefore for a particular policy μ

$$\log(T_\mu) \approx \frac{\alpha}{1 - \alpha} \log(\delta) - c_\mu$$

where c_μ depends on the policy and the system parameters (λ, \bar{s} , etc.) and α gives the distributional dependence of the policy. Thus, by plotting $\log(T_\mu)$ (or $\log(N_\mu)$) against $\log(\delta)$ and performing a linear regression, one can estimate α and hence the distributional dependence of the policy μ . We would expect a value of $\alpha = \frac{1}{2}$ for unbiased policies and a value of $\alpha = \frac{2}{3}$ for policies that behave like the optimal biased policy. Note that since $\log(\cdot)$ is increasing and $\log(\delta) < 0$, higher values of α imply lower system times.

Figure 2 shows a log-log plot of the sample average number in the system as a




Figure 2. Simulation results for SFC and NN policies for general demand distribution

function of δ for our simulation runs. The least squares estimate of the slope of each line is shown in Figure 2 as well. For the SFC policy, the estimated slope of 0.80 corresponds to $\alpha = 0.44$ while for the NN policy, the slope of 1.37 implies $\alpha = 0.58$. These values suggest that the SFC policy performs approximately like an unbiased policy since its value of α is close to $\frac{1}{2}$, though its performance appears to be somewhat worse (higher α) than a purely unbiased policy. The NN policy, on the other hand, seems to perform better than a purely unbiased policy but not as well as an optimal biased policy; that is, it achieves a higher value of α than an unbiased policy could, but does not achieve a value of $\alpha = \frac{2}{3}$ as in an optimally biased policy.

These results suggest a means of characterizing other analytically intractable policies; namely estimate γ_μ and α_μ as we did above and use the approximation

$$T_\mu \approx \gamma_\mu^2 \frac{\lambda \Xi(\alpha_\mu)}{m^2 v^2 (1 - \rho)^2}.$$

For example, this estimation might be performed using operating data from a ‘live’ system and the results used to evaluate current operating practice.

5. Relationship between biased and unbiased behavior

To review, we have determined that

$$T^* = \Theta \left(\frac{\lambda \Xi}{m^2 v^2 (1 - \rho)^2} \right)$$

where for the uniform demand case, $\Xi = A$, for the spatially unbiased general demand case $\Xi = (\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx)^2$ and for spatially biased general demand case $\Xi = (\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx)^3$. We next briefly examine the relationship among these various distributional behaviors.

Since unbiased service is a constraint, the system time of the optimal biased policy should be lower than the optimal unbiased policy for all densities f . This is indeed the case as shown by the following proposition, which also gives the relationship of the general distribution case to the uniform case.

Proposition 1. For any continuous density function $f(x)$ defined over the region \mathcal{A} of area A

$$A \cong \left[\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx \right]^2 \cong \left[\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx \right]^3$$

with equality holding throughout if and only if $f(x) = 1/A, \forall x \in \mathcal{A}$.

Proof. The proof requires the following inequality of Hardy et al. [9].

Lemma 6 (Hardy et al.). If $\alpha > 1$ or $\alpha < 0$, $g(x) \geq 0$ and $h(x) \geq 0$ then

$$\int g(x)^{1-\alpha} h(x)^\alpha dx \geq \left(\int g(x) dx \right)^{1-\alpha} \left(\int h(x) dx \right)^\alpha$$

with equality if and only if $g(x)/h(x)$ is constant for all x .

For the first inequality in our proposition, take $g(x) = f(x)$, $h(x) = f^{\frac{1}{2}}(x)$ and $\alpha = 2$ above, and note that $g(x)^{1-\alpha}h(x)^\alpha = f^{-1}(x)f(x) = 1$ which implies that $\int_{\mathcal{A}} g(x)^{1-\alpha}h(x)^\alpha dx = \int_{\mathcal{A}} dx = A$. Also,

$$\left(\int_{\mathcal{A}} g(x) dx\right)^{1-\alpha} \left(\int_{\mathcal{A}} h(x) dx\right)^\alpha = \left(\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx\right)^2.$$

Thus,

$$A \geq \left(\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx\right)^2$$

with equality if and only if $f(x)/f^{\frac{1}{2}}(x) = f^{\frac{1}{2}}(x)$ is constant for all x , which implies $f(x) = 1/A, \forall x \in \mathcal{A}$.

For the second inequality, take $g(x) = f^{\frac{2}{3}}(x)$, $h(x) = f^{\frac{1}{2}}(x)$ and $\alpha = -2$ above and note that $g(x)^{1-\alpha}h(x)^\alpha = f^2(x)f^{-1}(x) = f(x)$ and $\int_{\mathcal{A}} f(x) dx = 1$ we obtain

$$\left(\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx\right)^3 \left(\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx\right)^{-2} \leq 1.$$

Equality holds above if and only if $f^{\frac{2}{3}}(x)/f^{\frac{1}{2}}(x) = f^{\frac{1}{6}}(x)$ is constant for all x , again implying $f(x) = 1/A, \forall x \in \mathcal{A}$.

Remark. Proposition 1 says that a uniform density is the *worst possible* and that any deviation from uniformity in the demand distribution will strictly lower the optimal mean system time in either the unbiased or biased case. In addition, allowing biased service will result in a strict reduction of the optimal mean system time for any non-uniform distribution f . Also, note that when the density is uniform there is nothing to be gained by providing biased service.

One may question how different the system times for a biased and unbiased policy may be in general. That is, how much can one gain by discriminating according to location? Or, alternatively, how much does one lose by imposing an unbiased service constraint? The answer is that in the worst case the two can be arbitrarily far apart. This is illustrated by the simulation example in Figure 1. For the density used in this example, it is straightforward to show that for a fixed $\delta > 0$ and $\varepsilon \rightarrow 0$,

$$\left[\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx\right]^2 = \delta(1 - \varepsilon) + O(\varepsilon^{\frac{1}{2}})$$

and

$$\left[\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx\right]^3 = \delta^2(1 - \varepsilon) + O(\varepsilon^{\frac{1}{3}}).$$

Thus, there exists a constant c such that in heavy traffic

$$\frac{T_F^*}{T_D^*} \geq c \frac{\left[\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx\right]^2}{\left[\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx\right]^3} \rightarrow \frac{c}{\delta} \quad \text{as } \varepsilon \rightarrow 0,$$

where T_F^* and T_D^* are, respectively, the optimal unbiased and biased mean system times. Since $\delta > 0$ can be arbitrarily small, this says that in heavy traffic the cost of the optimal unbiased policy can be unbounded relative to the cost of the optimal biased policy.

Intuitively, one can explain the behavior of this example as follows. In an unbiased policy, the few points that fall in the large regions \mathcal{A}_2 must be visited as regularly as the large number of points that fall in the much smaller region \mathcal{A}_1 . However, visiting the points in \mathcal{A}_2 is time-consuming since they are typically far away from neighboring points. These infrequent yet time-consuming trips to demands in \mathcal{A}_2 impose large delays on the demands in \mathcal{A}_1 , which in turn drags down the overall mean system time. In a biased policy, we can allow the relatively small number of demands in \mathcal{A}_2 to wait much longer than the demands in \mathcal{A}_1 . The demands in \mathcal{A}_2 then build up and can be serviced more efficiently with larger tours. This frees up more vehicle time to service the much higher fraction of demands that land in \mathcal{A}_1 , improving their system time. The net result is to reduce the overall system time.

6. On the tightness of the lower bounds for the general case

In the proof of Lemma 4, one can see that very little of the vehicle routing structure inherent in the DTRP was used. Indeed, we only used the fact that the service was sequential (i.e. one demand served at a time) and the normalized waiting times satisfied P3. This allowed us to establish that the mean number left behind by a departure from any given region was asymptotically the same as the time-average number in queue in that region. The bound therefore applies to any system in which points arrive randomly to a Euclidean region and are then removed sequentially according to some given rule. For example, we might remove a point after it spends a constant amount of time τ in the system, in which case the expected nearest-neighbor distance $E[Z^*]$ and the mean number in queue N would also satisfy Lemma 4. A DTRP policy, in this sense, simply defines one such rule for removing points; namely, remove a point after a vehicle following a given policy μ has completed its on-site service. In this section, we show that the lower bound in Lemma 4 is in fact tight within this broader class of *removal rules*, and therefore more vehicle routing features of the DTRP need to be exploited if one wants to improve on these bounds.

6.1. An optimal removal rule. As in the DTRP, consider a region \mathcal{A} that receives arrivals according to a renewal process with intensity λ . The locations of arriving points are i.i.d. and distributed according to a general spatial density $f(x)$. Points are removed from the system according to the following rule.

Optimal removal rule. Each arrival of a new point triggers a *round* of removals. A round of removals proceeds as follows: the oldest point in the system that is within a radius z of any neighboring point is removed. ($z > 0$ is an arbitrarily small

constant.) The second oldest point with z of any of the remaining points is then removed, etc. The round continues until no more points are left within z of any other point. Though these removals are sequenced, we assume the round of removals takes place instantaneously. This process is repeated for every arriving point.

Theorem 6. For the optimal removal rule,

$$E[Z^*] = \frac{2}{3\sqrt{\pi}} \frac{1}{\sqrt{N}} \int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx.$$

Proof. We first analyze this policy for the uniform demand case. Note that at the end of a round, all points in the system are more than a distance z from their nearest neighbor. Also, arriving points are never eliminated in the round of removals that they initiate. This is because all points within a radius z of the arriving point are necessarily older and thus will be eliminated before the current arrival is considered. Similarly, all points in the system at the time of an arrival that are within a distance z of the arrivals location will be eliminated during its round because the arriving point is always the newest.

Given these observations, we see that a point waits in the system until a subsequent arrival falls within a distance z of it, at which point it is eliminated by the round of removals generated by this arrival. Since the probability that an arrival falls within z of any given location is $\pi z^2/A$ (ignoring edge effects because z is small) and the mean interarrival time of points is $1/\lambda$, the waiting time, W , under this policy is

$$W = \frac{A}{\lambda\pi z^2}.$$

We next determine the expected nearest-neighbor distance at the time of removal, $E[Z^*]$. Consider the removal epoch of a point i whose location we denote x_i . Note that at the removal epoch there is only one point within a radius z of x_i , namely the point that initiated the round of removals. Thus, the arriving point that triggers the removal of i is always the nearest neighbor to x_i . Since the arriving point's location is uniformly distributed within the circle of radius z about x_i , we have

$$\begin{aligned} E[Z^*] &= \int_0^z P\{Z^* > x\} dx \\ &= \int_0^z \left(1 - \frac{\pi x^2}{\pi z^2}\right) dx \\ &= \frac{2}{3}z. \end{aligned}$$

Using the expression for W above we have

$$z = \sqrt{\frac{A}{\lambda W \pi}} = \sqrt{\frac{A}{\pi N}},$$

which substituted into the expression for $E[Z^*]$ implies

$$E[Z^*] = \frac{2}{3\sqrt{\pi}} \sqrt{\frac{A}{N}}.$$

Comparing this to the bound in Lemma 4 and recalling that $\phi(x) = f(x) = 1/A$ for the spatially unbiased, uniform case proves the theorem for the uniform case.

This result can be extended to the non-uniform case by taking the radius z above to be a function of a point's location x ; that is, $z(x)$. Define

$$z(x) = \sqrt{\frac{\varepsilon}{f(x)\pi}},$$

where $\varepsilon > 0$ is an arbitrarily small constant. Note that the conditional wait given that a point arrives at location x satisfies (for sufficiently small $z(x)$)

$$E[W | X = x] = \frac{1}{\lambda f(x)\pi z^2(x)} = \frac{1}{\lambda \varepsilon}$$

and is therefore the same as the unconditional waiting time W . Using this observation, we can write $z(x)$ as follows:

$$z(x) = \sqrt{\frac{1}{Nf(x)\pi}}.$$

For the same reasons as in the uniform case,

$$E[Z^* | X = x] = \frac{2}{3}z(x) = \frac{2}{3\sqrt{\pi}} f^{-\frac{1}{2}}(x) N^{-\frac{1}{2}}.$$

Unconditioning implies

$$E[Z^*] = \frac{2}{3\sqrt{\pi}} \frac{1}{\sqrt{N}} \int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx,$$

which establishes the theorem for general $f(x)$ as well.

Remark. We have not discovered an analogous example for the general biased case.

Remark. The bound in Lemma 4 is in essence a dynamic counterpart to the following static nearest neighbor bound for n uniformly distributed points in a region of area A :

$$E[Z^*] \geq \frac{1}{2} \frac{\sqrt{A}}{\sqrt{n}},$$

which is used in the probabilistic analysis of Euclidean problems such as the TSP, matching and minimum spanning tree [15]. In the same sense that this nearest neighbor bound is weak for the static TSP, one can see that the bound of Lemma 4

is likely to be weak for the DTRP. This suggests that the provable performance bound of 1.83 for the unbiased and biased policies is indeed too pessimistic.

7. Further extensions

7.1. *General demand distributions and capacitated vehicles.* Most of the results for the general demand distributions extend to the capacitated vehicle case as well. The capacity constraint we consider is an upper bound of q on the number of demands a vehicle can serve before it must visit the depot at x_0 . We let $\bar{r} \equiv E[||X - x_0||]$ denote the average distance from a demand location to the depot. We shall only summarize results in this section since the analysis closely parallels the arguments we have seen in previous sections and in [6] and [7].

By using the more general bound on the nearest neighbor distance $E[Z^*]$ of Lemma 4 in the arguments of [7], one can show the following theorems.

Theorem 7. Within the class of spatially unbiased policies,

$$\lim_{\rho + 2\lambda\bar{r}/(mvq) \rightarrow 1} T^* \left(1 - \rho - \frac{2\lambda\bar{r}}{mvq}\right)^2 \geq \frac{\gamma^2 \lambda \left(1 + \frac{1}{q}\right)^2 \left[\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx\right]^2}{9 m^2 v^2}$$

where $\gamma \geq 2/3\sqrt{\pi}$.

Theorem 8. With the class of spatially biased policies

$$\lim_{\rho + 2\lambda\bar{r}/(mvq) \rightarrow 1} T^* \left(1 - \rho - \frac{2\lambda\bar{r}}{mvq}\right)^2 \geq \frac{\gamma^2 \lambda \left(1 + \frac{1}{q}\right)^2 \left[\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx\right]^3}{9 m^2 v^2}$$

where $\gamma \geq 2/3\sqrt{\pi}$.

A provably good unbiased policy for the finite-capacity case can be obtained by modifying the unbiased policy from Theorem 3 as follows: as sets of size n/k are formed, partition these sets into feasible tours of at most q points using the tour partitioning heuristic of Haimovich and Rinnooy Kan [8] as was done for the modified q TP policy in [7]. Serve these sets FCFS and optimize over n . For large k the resulting system time, T_{qU} , then satisfies

$$T_{qU} \sim \frac{\lambda \beta^2 \left(1 - \frac{1}{q}\right)^2 \left(\int_{\mathcal{A}} f^{\frac{1}{2}}(x) dx\right)^2}{2m^2 v^2 \left(1 - \rho - \frac{2\lambda\bar{r}}{mvq}\right)^2},$$

which implies the same performance guarantee as in the uniform case (cf. [7]).

An identical tour partitioning modification applied to the sets formed in the spatially biased policy of Theorem 5 gives a policy with a system time, T_{qB} ,

satisfying

$$T_{qB} \sim \frac{\lambda \beta^2 \left(1 - \frac{1}{q}\right)^2 \left(\int_{\mathcal{A}} f^{\frac{2}{3}}(x) dx\right)^3}{2m^2 v^2 \left(1 - \rho - \frac{2\lambda \bar{f}}{mvq}\right)^2}.$$

These policies and bounds describe the behavior of the most general version of the DTRP we have seen thus far and give a comprehensive picture of how a rich set of parameters influences congestion in dynamic vehicle routing systems.

7.2. *High dimensions.* Most of the DTRP bounds and policies can be extended to Euclidean subsets \mathcal{A} of \mathbb{R}^d for arbitrary dimension d . We examine this extension briefly in this section.

Modifications to the proofs of Theorems 1 and 2 give the following bounds.

Theorem 9. Within the class of spatially unbiased policies,

$$\lim_{\rho \rightarrow 1} T^*(1 - \rho)^d \geq \frac{\gamma(d)^d \lambda^{d-1} \left[\int_{\mathcal{A}} f^{(d-1)/d}(x) dx\right]^d}{m^d v^d}$$

where

$$\gamma(d) = \frac{d}{d+1} \left(\frac{d}{d+1}\right)^{1/d} \left(\frac{1}{c_d}\right)^{1/d}$$

and

$$c_d = \frac{\pi^{d/2}}{\Gamma\left(\frac{d}{2} + 1\right)}$$

is the volume of the ball of unit radius in dimension d ($\Gamma(x)$ is the usual gamma function).

Theorem 10. Within the class of spatially biased policies

$$\lim_{\rho \rightarrow 1} T^*(1 - \rho)^d \geq \frac{\gamma(d)^d \lambda^{d-1} \left[\int_{\mathcal{A}} f^{d/(d+1)}(x) dx\right]^{d+1}}{m^d v^d}$$

where

$$\gamma(d) = \frac{d}{d+1} \left(\frac{1}{d+1}\right)^{1/d} \left(\frac{1}{c_d}\right)^{1/d}.$$

Again, similar results holds for the capacitated problem, in which case $(1 - \rho)$ becomes $(1 - \rho - (2\lambda\bar{r}/mvq))$ in the above bound and also $\gamma(d)$ is replaced by $\gamma(d)/3$.

In a similar manner, one can analyze the various service policies in d dimensions. The results parallel those in the two-dimensional case; namely, there are constants $\gamma_\mu(d)$ that depend only on the policy and the dimension d such that the system time, T_μ , satisfies

$$T_\mu \sim \gamma_\mu^d(d) \frac{\lambda^{d-1} \Xi(d)}{m^d v^d (1 - \rho)^d} \text{ as } \rho \rightarrow 1$$

where $\Xi(d) = V$ for the uniform case,

$$\Xi(d) = \left[\int_{\mathcal{A}} f^{(d-1)/d}(x) dx \right]^d$$

for the spatially unbiased case and

$$\Xi d = \left[\int_{\mathcal{A}} f^{d/(d+1)}(x) dx \right]^{d+1}$$

for the spatially biased case. For example, the modified TSP policy in d dimensions has a constant value of $\beta(d)/2^{1/d}$, where $\beta(d)$ is the d -dimensional TSP constant.

An interesting result is found by examining this policy for $d \rightarrow \infty$. In [5], it was conjectured and subsequently proved in [18] that for $d \rightarrow \infty$, $\beta(d) \sim \sqrt{d}/\sqrt{2\pi e}$. By using the fact that for $d \rightarrow \infty$, $d/(d+1) \sim 1$, $(1/(d+1))^{1/d} \sim 1$ and

$$\Gamma(\frac{1}{2}d + 1) \sim \sqrt{2\pi} \left(\frac{d}{2}\right)^{d/2 + \frac{1}{2}} \exp\left(-\frac{d}{2}\right),$$

it is straightforward to show that

$$\gamma(d) \sim \frac{\sqrt{d}}{\sqrt{2\pi e}}$$

as $d \rightarrow \infty$ as well. Therefore we have the following theorem.

Theorem 11. For the uncapacitated, m -server DTRP, the modified TSP policy is an optimal heavy traffic policy asymptotically as $d \rightarrow \infty$.

This theorem gives further evidence of the asymptotic optimality of the modified TSP policy, which we have conjectured is optimal for $d = 2$.

8. Conclusions

We analyzed dynamic vehicle routing problems in Euclidean regions under general distributional assumptions. The analysis yields simple expressions for the

system time that provide structural insight into the effects of traffic intensity, on-site service characteristics, the number, speed and capacity of vehicles employed, service region size, the distribution of demand locations and bias constraints. Such insights can be used to develop strategic planning models for terminal location, fleet sizing and districting. We see such strategic planning models as a potentially fruitful area for further applied research.

A recurring finding in our analysis is that static vehicle routing methods when properly adapted can yield near optimal or perhaps even optimal policies for dynamic routing problems. This is an encouraging result on several levels. On a theoretical level, it suggests that there is indeed a connection between the static and dynamic problems; that is, the DTRP has geometrical characteristics that are intimately related to the corresponding characteristics for static VRPs. On a practical level, the results imply that most of the exact algorithms, heuristics and insights which have been developed over the years of investigation of static VRPs are not irrelevant in this context and can in fact form the basis for effective policies in dynamic, stochastic environments.

References

- [1] BARTHOLDI, J. J. AND PLATZMAN, L. K. (1982) An $O(N \log N)$ planar traveling salesman heuristic based on space filling curves. *Operat. Res. Lett.* **1**, 121–125.
- [2] BARTHOLDI, J. J. AND PLATZMAN, L. K. (1988) Heuristics based on space filling curves for combinatorial problems in Euclidean space. *Management Sci.* **34**, 291–305.
- [3] BEARDWOOD, J., HALTON, J. AND HAMMERSLEY, J. (1959) The shortest path through many points. *Proc. Camb. Phil. Soc.* **55**, 299–327.
- [4] BERTSIMAS, D. AND NAKAZATO, D. (1991) The general distributional Little's law and its applications. *Operat. Res.* To appear.
- [5] BERTSIMAS, D. AND VAN RYZIN, G. (1990) An asymptotic determination of the minimum spanning tree and minimum matching constants in geometrical probability. *Oper. Res. Lett.* **9**, 223–231.
- [6] BERTSIMAS, D. AND VAN RYZIN, G. (1990) A stochastic and dynamic vehicle routing problem in the euclidean plane. *Operat. Res.* **39**, 601–615.
- [7] BERTSIMAS, D. AND VAN RYZIN, G. (1993) Stochastic and dynamic vehicle routing in the euclidean plane with multiple capacitated vehicles. *Operat. Res.* **41**, 60–76.
- [8] HAIMOVICH, M. AND RINNOOY KAN, A. H. G. (1985) Bounds and heuristics for capacitated routing problems. *Math. Operat. Res.* **10**, 527–542.
- [9] HARDY, G. H., LITTLEWOOD, J. E. AND PÓLYA, G. (1934) *Inequalities*. Cambridge University Press.
- [10] IGLEHART, D. L. AND WHITT, W. (1970) Multiple channel queues in heavy traffic, I and II. *Adv. App. Prob.* **2**, 150–177 and 355–364.
- [11] JOHNSON, D. (1988) Talk presented at the Mathematical Programming Symposium, Tokyo.
- [12] KARP, R. (1977) Probabilistic analysis of partitioning algorithms for the traveling salesman in the plane. *Math. Operat. Res.* **2**, 209–224.
- [13] KARP, R. M. AND STEELE, J. M. (1985) Probabilistic analysis of heuristics. In *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*, ed. E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan and D. B. Shmoys. Wiley, Chichester.
- [14] LAWLER, E. L., LENSTRA, J. K., RINNOOY KAN, A. H. G. AND SHMOYS, D. B. (EDS.) (1985) *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*. Wiley, Chichester.
- [15] PAPADIMITRIOU, C. H. (1978) The probabilistic analysis of matching heuristics. *Proc. 15th Annual Allerton Conf. on Communication, Control and Computing*, pp. 368–378.

- [16] PLATZMAN, L. K. AND BARTHOLDI, J. J. (1983) Spacefilling curves and the planar traveling salesman problem. PDRC Technical Report 83-02, Georgia Institute of Technology.
- [17] PSARAFTIS, H. (1988) Dynamic vehicle routing problems. In *Vehicle Routing: Methods and Studies*, ed. B. Golden and A. Assad, North-Holland, Amsterdam.
- [18] RHEE, W. (1991) On the TSP in many dimensions. *Random Structures*. To appear.
- [19] STEELE, J. M. (1981) Subadditive euclidean functionals and nonlinear growth. *Ann. Prob.* **9**, 365–376.
- [20] WOLFF, R. W. (1982) Poisson arrivals see time averages. *Operat. Res.* **30**, 223–231.