Stochastic and Dynamic Vehicle Routing with General Demand and Interarrival Time Distributions

Dimitris J. Bertsimas
Garrett van Ryzin

WP# 3314-MS

June 24, 1991

Stochastic and Dynamic Vehicle Routing with General Demand and Interarrival Time Distributions

Dimitris J. Bertsimas * Garrett van Ryzin †‡

June 24, 1991

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 analysis ([5] and [6]) on this problem, we needed to assume 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 distribution and arrivals follow a general renewal process. Further, we improve significantly the best known lower bounds for this class of problems and construct policies which 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 which offers important structural insight to 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.

Key words. dynamic vehicle routing, general demand distribution, general interarrival times, bounds, heuristics.

---

*Dimitris Bertsimas, Sloan School of Management, MIT, Rm E53-359, Cambridge, MA 02139.

†Garrett van Ryzin, Operations Research Center, MIT, Cambridge, MA 02139.

‡The research of both authors was partially supported by the National Science Foundation under grant DDM-9014751 and by grants from Draper Laboratories and the UPS Foundation.

1 Introduction

Dynamic vehicle routing problems occur when one has to visit customers (*demands*) that arrive sequentially over a period of time. The objective is to schedule these visits in a way that is economical yet also provides an acceptable service level (wait for delivery/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 prior to 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 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/service.

There are many practical settings in which such problems arise. Any distribution system which receives orders in realtime and makes deliveries based on these orders (courier services, deliveries of flowers and 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/highway maintenance. A third important application is finished goods distribution from factories to retailers. Arrivals in this context correspond to the completion of a unit at the factory that is designated for a particular retailer. The average waiting time in this case not only represents a service level to the retailer, but, for a fixed production rate, is also proportional to the average inventory held in the distribution channel. A specific example of this type is the delivery of automobiles from assembly plants to local dealers. We refer the reader to the excellent general discussion in [15] and to the more specific discussions in [5] and [6] for more background on dynamic vehicle routing problems,.

1.1 The Dynamic Traveling Repairman Problem (DTRP)

In [5] and [6] 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 random variables that are independent and uniformly distributed throughout the service region. (These locations are random only in the stochastic processes sense mentioned above; namely that at time t we know with *certainty* the locations of demands that arrived prior to time t, but future demand locations form and 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 [6], has a guarantee of $T_\mu/T^* \leq 3.6$. In [6] we also extended the model 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 according to an arbitrary, continuous density defined over the service region and arrivals are generated according to a general renewal process. (Formal definitions are given in Section 2.) These extensions are important for both practical and theoretical reasons. On the practical side, one would clearly like to relax the assumption of uniformly distributed demand locations, since this distribution is rarely encountered in real systems. Hence, to apply such models, an understanding of the problem cost and solution structure under nonuniform distributions is essential. Also, while the Poisson arrival assumption may be appropriate for certain repair systems, other applications (*e.g.* finished goods distribution from assembly plants) have less variable interarrival times: therefore, it is important to understand to what extent the results hold for more regular and/or variable arrival sequences.

From a theoretical perspective, these generalizations are challenging because uniformity and PASTA [17] were exploited heavily in our previous analysis; thus, entirely different analytical techniques are required. These techniques are introduced and developed in Section 3. Our new analytical approach has the added benefit of strengthening the lower bound on T^* by a factor of two, and 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 ratio is in fact one, and thus 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 which 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 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, improvements can only be made by exploiting more 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 Problem Definition

A total of m identical vehicles operate in a bounded service region $\mathcal{A} \subset \Re^2$. Vehicles travel at a constant, finite velocity v. Demands arrive to $\mathcal{A}$ according to a general renewal process with intensity λ. The interarrival times have finite variance σ_a^2 and Laplace transform $A^*(s)$. At each location, vehicles spend some time s in on-site service that is i.i.d. and generally distributed with finite first and second moments denoted by $\overline{s}$ and $\overline{s^2}$ respectively.

Demand locations $\{X_i;\ i \geq 1\}$ are i.i.d. and distributed according to a continuous density $f(x)$ defined over $\mathcal{A}$. They become known (are realized) at a demand's arrival epoch. The notation $f(x)$ is short for $f(x_1, x_2)$ ($x = [x_1\ x_2]$). Likewise, we write $\int f(x)dx$ for $\int\int f(x_1, x_2)dx_1dx_2$. The density $f(x)$ satisfies

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

and

$$\int_{\mathcal{A}} f(x)dx = 1.$$

For our analysis, we require the technical condition that $0 < \underline{f} \leq f(x) \leq \overline{f} < \infty, \forall\ x \in \mathcal{A}$.

A policy for routing the vehicles is called *stable* if the number of demands in the system is bounded almost surely for all times t. If a policy is stable, $\rho \equiv \frac{\lambda \overline{s}}{m}$ is the fraction of total vehicle time spent in on-site service. The term *heavy traffic* is used to denote the condition $\rho \longrightarrow 1$. The mean waiting time in this system is denoted generically by W and the mean system time (wait in queue plus on-site service) by $T = W + \bar{s}$.

Let $\mathcal{M}$ denote the subset of all stable, stationary policies μ in which decisions are taken only at service completion epochs and consist of choosing one of the demands in the system to visit next or alternatively choosing to visit a fixed depot location x_0. We point out that in the uniform demand case discussed in [5] and [6], we allow for a slightly more general class of policies in which vehicles can wait at any location x and change destinations at any point in time. Note that within the class $\mathcal{M}$ a vehicle can choose the option of staying put only at the depot x_0. (At the depot no restrictions on the timing of decisions is imposed.) This turns out to be an important property for our subsequent analysis.

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

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

We let T^* denote the infimum on the right hand side above. A policy for which $\frac{T_\mu}{T^*}$ is bounded is said to have a *constant factor guarantee.* If $\lim_{\rho \rightarrow 1} \frac{T_\mu}{T^*}$ is bounded then the policy μ is said to have a constant factor guarantee in heavy traffic.

We shall also need two definitions mentioned informally in Section 1. In these definitions, X is the location of a randomly chosen demand and W is its waiting time.

Definition 1 *A policy μ is called spatially unbiased if*

$$E[W|X \in \mathcal{S}] = W \quad \forall\ \mathcal{S} \subseteq \mathcal{A}.$$

and

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

$$E[W|X \in \mathcal{S}_1] > E[W|X \in \mathcal{S}_2],$$

If we let $W(x) \equiv E[W|X = x]$, then observe that for spatially unbiased policies $W = W(x)$ for all x while for spatially biased policy W is given only by the more general relation

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

3 Heavy Traffic Lower Bounds

In this section, we derive our main lower bounds on T^*. The bounds are established using a different proof technique than that used in [5] and [6] that not only allows us to consider general spatial distributions and arrival processes, but also improves on the constant value for the Poisson, uniform case ($f(x) = 1/A$). We begin by proving some lemmas related to "spatial queues" in the system that culminate in a generic lower bound on T^*. This bound is then specialized to the spatially biased and unbiased cases to arrive at our main lower bounds.

3.1 Preliminary Lemmas

3.1.1 Spatial Queues on Subsets of $\Re^2$

We associate a queue with every $\mathcal{S} \subseteq \Re^2$, henceforth referred to as simply as the *the queue* $\mathcal{S}$, 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 $N(\mathcal{S})$ denote the time average number of customers 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 $N = N(\mathcal{A})$.

Let $\mathcal{C}(x, z)$ be the set of points within a distance of z from the location x (*i.e.* $\mathcal{C}(x, z) = \{y|\ \|y - x\| \leq z\}$). For all $x \in \mathcal{A}$, we define the following limit, which is essentially the time average density of demands in queue at locations x:

$$\phi(x) = \frac{1}{N} \lim_{z \to 0} \frac{N(\mathcal{C}(x, z))}{\pi z^2}. \qquad (1)$$

We shall assume that if the system is stable, then this limit exists. Further, we will need a regularity condition on ϕ; namely, that if $\underline{f} \leq f(x) \leq \overline{f}$ for all $x \in \mathcal{A}$, then $\underline{\phi} \leq \phi(x) \leq \overline{\phi}$ for all $x \in \mathcal{A}$, where these bounds on $\phi(x)$ may depend on $f(x)$ but not on N. Intuitively, this condition says that the density of demands in queue at any location x in the system grows uniformly as $N \rightarrow \infty$.

From the definition of $\phi(x)$ and the linearity of expectation, we have for any subset $\mathcal{S}$ of $\Re^2$ that

$$N(\mathcal{S}) = N \int_{\mathcal{S}} \phi(x)dx. \tag{2}$$

Also, since $N(\mathcal{A}) = N$ and $N(\cdot)$ is always positive, $\phi(x)$ satisfies

$$\begin{aligned} \int_{\mathcal{A}} \phi(x)dx &= 1 && (3) \\ \phi(x) &\geq 0 \quad \forall \;\; x \in \mathcal{A}. && (4) \end{aligned}$$

It is these last two properties of ϕ that give it its interpretation as a density function.

3.1.2 The Arrival Process to "'Small" Subsets $\mathcal{S}$

We next show that if $\mathcal{S}$ is small, the expected number of demands left behind by a random departure from $\mathcal{S}$ is approximately the time average number in queue $N(\mathcal{S})$. This result which will be useful in our subsequent analysis.

Let the random variable $Y(\mathcal{S})$ denote an interarrival time of the queue $\mathcal{S}$, $p(\mathcal{S}) = \int_{\mathcal{S}} f(x)dx$ denote the probability that an arrival falls in the set $\mathcal{S}$ and $\lambda(\mathcal{S}) = \lambda p(\mathcal{S})$ denote the arrival rate to $\mathcal{S}$. Note that $Y(\mathcal{S})$ is a geometric sum of interarrival times, and thus its transform $F^*_{Y(\mathcal{S})}(s)$ is given by

$$F^*_{Y(\mathcal{S})}(s) = \frac{A^*(s)p(\mathcal{S})}{1 - A^*(s)(1 - p(\mathcal{S}))}. \tag{5}$$

Finally, let $n^+(\mathcal{S})$, a random variable, denote the number of customers left behind by a random departure from $\mathcal{S}$, $W(\mathcal{S})$ denote the waiting time in this queue and recall $N(\mathcal{S})$ denotes the time average number of demands in $\mathcal{S}$. These definitions allow us to state the following lemma:

Lemma 1 *Let $||\mathcal{S}||$ denote the area of $\mathcal{S}$, then*

$$E[n^+(\mathcal{S})] = N(\mathcal{S}) + o(||\mathcal{S}||).$$

In particular, if $\mathcal{S} = \mathcal{C}(x,z)$, then

$$\lim_{z \to 0} \frac{E[n^+(\mathcal{C}(x,z))]}{\pi z^2} = N\phi(x).$$

Proof

Note that if the region $\mathcal{A}$ *itself* has Poisson arrivals, then for all $\mathcal{S}$, $E[n^+(\mathcal{S})] = N(\mathcal{S})$. This follows from PASTA [17] and the fact that customers are served sequentially (one at a time). To prove the lemma, it is therefore sufficient to show that the normalized interarrival time

$$\hat{Y}(\mathcal{S}) \equiv \frac{Y(\mathcal{S})}{\lambda^{-1}(\mathcal{S})}$$

has an exponential distribution for $||\mathcal{S}|| \to 0$ and then invoke the above PASTA result.

Letting $F^*_{\hat{Y}(\mathcal{S})}(s)$ denote the transform of $\hat{Y}(\mathcal{S})$ and suppressing the argument $\mathcal{S}$ in $p(\mathcal{S})$ for brevity, we obtain

$$\begin{aligned} F^*_{\hat{Y}(\mathcal{S})}(s) &= \int_0^\infty e^{-st} dF_{\hat{Y}(\mathcal{S})}(t) \\ &= \int_0^\infty e^{-st} dF_{Y(\mathcal{S})}(\lambda^{-1}(\mathcal{S})t) \\ &= \int_0^\infty e^{-sp\lambda t} dF_{Y(\mathcal{S})}(t) \\ &= F^*_{Y(\mathcal{S})}(sp\lambda). \end{aligned}$$

Note that $||\mathcal{S}|| \to 0$ implies $p(\mathcal{S}) \to 0$. Therefore using (5), taking the limit as $p \to 0$ and applying L'Hospital's rule we obtain

$$\begin{aligned} \lim_{p \to 0} F^*_{\hat{Y}(\mathcal{S})}(s) &= \lim_{p \to 0} \frac{A^*(sp\lambda)p}{1 - A^*(sp\lambda)(1-p)} \\ &= \lim_{p \to 0} \frac{A^*(sp\lambda) + pA'(sp\lambda)s\lambda}{A^*(sp\lambda) - (1-p)A'(sp\lambda)s\lambda} \\ &= \frac{1}{1+s}, \end{aligned}$$

which is the transform of a exponential random variable with unit intensity.

□ (Lemma 1)

The key insight shown by this lemma is that sampling a renewal process with low probability generates a Poisson process. Thus, for small regions $\mathcal{S}$, the arrival process to the queue $\mathcal{S}$ is approximately Poisson.

3.1.3 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 (which ever is smaller) at the completion epoch of demand i (the "nearest neighbor" distance). That is, Z_i^* is the decision that minimizes d_i; thus, $E[Z_i^*] \leq E[d_i]$ and

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

We are now ready to state and prove the following key lemma:

Lemma 2

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

<u>Proof</u>

Consider a randomly tagged demand arriving at location X and condition on the event $\{X = x\}$. Recall that $n^+(\mathcal{C}(x,z))$ denotes the number of customers in the set $\mathcal{C}(x,z)$ at the completion epoch of this customer. Then,

$$P(Z^* \leq z | X = x) = P(n^+(\mathcal{C}(x,z)) > 0) \leq E[n^+(\mathcal{C}(x,z))], \tag{6}$$

where the last inequality is due to the fact that $n^+(\mathcal{C}(x,z))$ is a nonnegative, integer-valued random variable. Note that we have implicitly assumed that the depot (location x_0) is not within a radius z of x, else the probability above would be one. We consider this alternate case below.

Considering the service completion of our tagged demand as a departure from the queue $\mathcal{C}(x,z)$, we therefore have by Lemma 1 above that as $z \longrightarrow 0$,

$$E[n^+(\mathcal{C}(x,z))] = N(\mathcal{C}(x,z)) + o(z^2) = N \int_{\mathcal{C}(x,z)} \phi(x)dx + o(z^2).$$

Expressing the integral above in terms of its asymptotic ($z \longrightarrow 0$) value and substituting into the bound (6) implies

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

Defining $c \equiv N\pi\phi(x)$, we therefore have

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

As mentioned, this bound is valid as long as the depot at x_0 is not within a radius $c^{-1/2} = [N\pi\phi(x)]^{-1/2}$ of the location x. Let $\mathcal{D}(N) = \{x|\ ||x - x_0|| \leq [N\pi\phi(x)]^{-1/2}\}$ denote the set of points for which the bound is not valid. We next establish that for large N, the contribution to the lower bound from the set $\mathcal{D}(N)$ is negligible – and it here that we need our technical assumptions on $f(x)$ and $\phi(x)$. First note that $\phi(x) \geq \underline{\phi}\ \ \forall N$ implies $\int_{\mathcal{D}(N)} dx \leq O(1/N)$. Using the trivial bound $P(Z^* > z|X = x) \geq 0$ for the points in $\mathcal{D}(N)$ and removing the conditioning $\{X = x\}$ implies,

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

$$= \frac{2}{3\sqrt{\pi N}} \int_{\mathcal{A}} \phi^{-1/2}(x) f(x) dx - o(N^{-1/2}),$$

which shows the contribution due to $\mathcal{D}(N)$ is indeed insignificant. Multiplying both sides above by $\sqrt{N}$ and taking the limit as $N \to \infty$ then proves the lemma.

□ (Lemma 2)

Lemma 2 can be used to prove the following intermediate bound on the optimal system time T^*:

Lemma 3 *There exists a constant γ such that*

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

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

Proof

Consider the following necessary condition for stability

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

Using the fact that $E[Z^*] \leq \bar{d}$, multiplying the second term on the left hand side above by $\frac{\sqrt{N}}{\sqrt{N}}$ 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 $\rho \to 1$, we must have $N \to \infty$. Therefore taking the limit as $\rho \to 1$ (and consequently $N \to \infty$) on both sides above and applying Lemma 2 we obtain

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

where $\gamma \geq \frac{2}{3\sqrt{\pi}}$. Squaring both sides and using $T \geq W = \frac{N}{\lambda}$ we obtain Lemma 3.

□ (Lemma 3)

3.2 A Spatially Unbiased Lower Bound

As mentioned, Lemma 3 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*

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

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

Before proving this theorem, we note that it differs from the heavy traffic bound in [5] and [6] in that it is an asymptotic bound while our earlier bounds are valid for all values of ρ; however, Theorem 1 improves on the constant value γ by a factor of $\sqrt{2}$ and thus increases the bound by a factor of two.

Proof

Consider the queue $\mathcal{C}(x,z)$ and recall that $W(\mathcal{C}(x,z))$ and $N(\mathcal{C}(x,z))$ are the mean wait and mean number within subset $\mathcal{C}(x,z)$ respectively. By Little's Theorem,

$$N(\mathcal{C}(x,z)) = (\lambda \int_{\mathcal{C}(x,z)} f(\xi)d\xi)\ W(\mathcal{C}(x,z)).$$

However, if the policy being used is spatially unbiased this implies that $W(\mathcal{C}(x,z)) = W$. Substituting this above and recalling that $N(\mathcal{C}(x,z)) = N\int_{\mathcal{C}(x,z)} \phi(\xi)d\xi$ we obtain

$$N \int_{\mathcal{C}(x,z)} \phi(\xi)d\xi = \lambda W \int_{\mathcal{C}(x,z)} f(\xi)d\xi,$$

which, since $N = \lambda W$, implies

$$\int_{\mathcal{C}(x,z)} \phi(\xi)d\xi = \int_{\mathcal{C}(x,z)} f(\xi)d\xi.$$

Letting $z \to 0$ above and noting that the above equality is true for all sets $\mathcal{C}(x,z)$ implies that $\phi(x) = f(x)\ \ \forall\ x \in \mathcal{A}$. Making the substitution $\phi(x) = f(x)$ in the bound in Lemma 3 we obtain the theorem.

□ (Theorem 1)

Note that $\phi(x) = f(x)$ also implies $\underline{\phi} = \underline{f}$ and $\bar{\phi} = \bar{f}$, which confirms our regularity assumption on $\phi(x)$.

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*

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

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

Proof

Since no assumption of unbiased service is made, consider the following minimization problem for the integral term in Lemma 3.

$$\begin{aligned} z^* &= \min \int_{\mathcal{A}} \phi^{-1/2}(x) f(x) dx \qquad (7) \\ \text{subject to} \quad & \int_{\mathcal{A}} \phi(x) dx = 1 \\ & \phi(x) \geq 0. \end{aligned}$$

Using the value z^* as a lower bound on the integral term in Lemma 3 will give us Theorem 1.

Note that the objective function is convex in $\phi(x)$ and the constraints are linear; thus, (7) is a convex program. Relaxing the equality constraint above with a multiplier, we obtain the following Lagrangian dual

$$\begin{aligned} z^*(\mu) &= \min_{\phi(x) \geq 0} \int_{\mathcal{A}} \phi^{-1/2}(x) f(x) dx + \mu \left[\int_{\mathcal{A}} \phi(x) dx - 1\right] \\ &= \int_{\mathcal{A}} \min_{\phi(x) \geq 0} \left[\phi^{-1/2}(x) f(x) + \mu \phi(x)\right] dx - \mu \qquad (8) \end{aligned}$$

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

$$-\frac{1}{2}[\phi^*(x)]^{-3/2} f(x) + \mu^* = 0 \quad \forall\, x \in \mathcal{A} \tag{9}$$

$$\int_{\mathcal{A}} \phi^*(x) dx = 1 \tag{10}$$

$$\phi^*(x) \geq 0 \tag{11}$$

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

$$\phi^*(x) = \left[\int_{\mathcal{A}} f^{2/3}(x) dx\right]^{-1} f^{2/3}(x) \tag{12}$$

$$\mu^* = \frac{1}{2}\left[\int_{\mathcal{A}} f^{2/3}(x) dx\right]^{3/2} \tag{13}$$

is such a pair. The fact that (7) is a convex program implies that these conditions are also sufficient to assure global optimality. Substituting the value $\phi^*(x)$ above into Lemma 3 gives us the theorem.

□ (Theorem 2)

Again, note that the regularity condition $\underline{\phi} \leq \phi^*(x) \leq \overline{\phi}$ is satisfied if $\underline{f} \leq f(x) \leq \overline{f}$.

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 [5] and [6].

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

$\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_k$ such that $\int_{\mathcal{A}_i} f(x)dx = \frac{1}{k}$ $i = 1, 2, \ldots, k$. (One could do this by "sweeping" the region from the depot using an arbitrary starting ray until $\int_{\mathcal{A}_1} f(x)dx = \frac{1}{k}$, continuing the sweep until $\int_{\mathcal{A}_2} f(x)dx = \frac{1}{k}$, 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 forming a TSP on the set and following it in an arbitrary directions. Optimize over n.

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

Proposition 1 *Let T_U^* be the optimal system time over the class of spatially unbiased policies. Then*

$$\frac{T_U}{T_U^*} \le \frac{\beta^2}{2\gamma^2} \approx 1.8 \qquad \text{as } \rho \to 1.$$

where $\beta \approx 0.72$ is the TSP constant in the Euclidean plane (see [3] and [10]).

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], we can therefore assert that almost surely

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

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

$$\begin{aligned}\frac{E[\tau]}{(n/k)} &= \bar{s} + \frac{1}{v\sqrt{n}}\sum_{i=1}^{k}\frac{E[L_i]}{\sqrt{n}} \\ &\sim \bar{s} + \frac{1}{v\sqrt{n}}\sum_{i=1}^{k}\beta\int_{\mathcal{A}_i} f^{1/2}(x)dx \\ &= \bar{s} + \frac{\beta}{v\sqrt{n}}\int_{\mathcal{A}} f^{1/2}(x)dx \qquad (14)\end{aligned}$$

To determine σ_T^2, consider the random variable L which is an equiprobable selection from the set of random variables $\{L_1, \ldots, L_k\}$. That is, L is the random variable such that $\sigma_T^2 = \frac{n}{k}\sigma_s^2 + \frac{1}{v^2}Var[L]$. Note by the above asymptotic behavior of $L_i/\sqrt{n}$ that for large n the random variable $L/\sqrt{n}$ approaches an equiprobable selection from the set of *constants* $\{\beta \int_{\mathcal{A}_1} f^{1/2}(x)dx, \ldots, \beta \int_{\mathcal{A}_k} f^{1/2(x)}dx\}$, and thus it follows that for fixed k

$$\frac{Var[L]}{n} = Var[\frac{L}{\sqrt{n}}] = O(1).$$

Hence

$$\frac{\sigma_T^2}{(n/k)} = \sigma_s^2 + O(1). \tag{15}$$

We will use (14) and (15) 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 $\sum$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 Inglehart and Whitt:

Theorem 3 (Inglehart and Whitt [9]) *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, $i = 1, 2, \ldots, k$. Let $1/\mu_j$ and $\sigma_{b_j}^2$ denote the mean and variance, respectively, of the service times at server $j = 1, 2, \ldots, m$. Define $\lambda \equiv \sum_{i=1}^{k} \lambda_i$ $\mu \equiv \sum_{j=1}^{m} \mu_j$ and $\rho \equiv \frac{\lambda}{\mu}$. Then as $\rho \to 1$ the mean waiting time in queue, W, satisfies*

$$W \sim \frac{\sum_{i=1}^{k} \lambda_i^3 \sigma_{a_i}^2 + \sum_{j=1}^{m} \mu_j^3 \sigma_{b_j}^2}{2\mu^2(1-\rho)}. \tag{16}$$

Let $\hat{\lambda}_i = \lambda/n$ denote the arrival rate of sets to region i and $\hat{\lambda} = \sum_{i=1}^{k} \hat{\lambda}_i = \frac{k\lambda}{n}$ 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

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

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

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

We shall use the fact that for large values of k, the right hand side above is approximately $\frac{1}{\lambda}\hat{\lambda}^2$, and in heavy traffic, $\hat{\lambda} \approx \frac{m}{E[\tau]}$. Using these facts and applying (14) and (15), we can therefore establish the following limit for W_{set}, the time a set waits in queue:

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

Note that the stability condition for this queue implies that $\rho + \frac{\lambda\beta}{mv\sqrt{n}}\int_{\mathcal{A}} f^{1/2}(x)dx < 1$, which implies

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

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

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

and

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

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

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

Making a change of variable to

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

the bound can be written

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

An approximate optimal value for y is

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

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

$$T_U \sim \frac{\lambda\beta^2(\int_A f^{1/2}(x)dx)^2(1+\frac{m}{k})}{2m^2v^2(1-\rho)^2},$$

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

□ (Proposition 1)

4.2 A Provably Good Spatially Biased Policy for Piece-Wise Uniform Demand

We next propose a policy that achieves a performance guarantee of $\frac{\beta^2}{2\gamma^2} \approx 1.8$ 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 $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_J$ such that $f(x) = \mu_j \ \ \forall\ x \in \mathcal{A}_j,\ j = 1, 2, \ldots, J$. In particular, for such a density note that

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

Though such densities is not perfectly general, one could approximate a continuous density by a piecewise continuous density and let the approximation become finer and finer to handle more general cases. Moreover, in practice a piecewise uniform density is probably adequate. The policy is defined as follows:

The Biased (B) TSP Policy Let $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_J$ be a partition of $\mathcal{A}$ such that $f(x) = \mu_j \ \ \forall\ x \in \mathcal{A}_j,\ j = 1, 2, \ldots, J$. Let A_j denote the area of $\mathcal{A}_j$. For a given positive integer k, partition each subset $\mathcal{A}_j$ further into $k_j = \mu_j^{2/3} A_j k$ regions of area $A_j/k_j = (\mu_j^{2/3} k)^{-1}$ (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 $\mu_j^{2/3} A_j$). 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) form 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 the system time of this policy be denoted T_B. We shall prove the following proposition:

Proposition 2 *If f is a piecewise uniform density and T_B^* is the optimal system time over the class of biased policies, then*

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

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 $\rho \longrightarrow 1$, $n \longrightarrow \infty$; therefore

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

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

$$\begin{aligned} \frac{E[\tau]}{(n/k)} &\longrightarrow \bar{s} + \frac{\beta}{v\sqrt{n}}\sum_{j=1}^{J} p_j\mu_j^{-1/3} \\ &= \bar{s} + \frac{\beta}{v\sqrt{n}}\sum_{j=1}^{J} \mu_j^{2/3}A_j \end{aligned} \quad (17)$$

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

$$\sigma_\tau^2 = \frac{n}{k}\sigma_s^2 + Var[L].$$

For large n, the random variable $\frac{kL}{\sqrt{n}}$ tends to a selection of constants from the set $\{\beta\mu_j^{-1/3}\}$ with probability p_j, and thus it follows that as $n \longrightarrow \infty$

$$Var[\frac{kL}{\sqrt{n}}] = k\frac{Var[L]}{(n/k)} = O(1)$$

and hence $\frac{Var[L]}{(n/k)} = \frac{O(1)}{k}$. Thus, for large n

$$\frac{\sigma_\tau^2}{n/k} = \sigma_s^2 + \frac{O(1)}{k}. \quad (18)$$

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

$$\begin{aligned} W^- &\leq \sum_{j=1}^{J} p_j \frac{1}{2}\left(\frac{n/k}{(p_j\lambda)/k_j}\right) \\ &= \frac{n}{2\lambda}\sum_{j=1}^{J}\frac{k_j}{k} \\ &= \frac{n}{2\lambda}\sum_{j=1}^{J}\mu_j^{2/3}A_j \end{aligned} \tag{19}$$

and

$$\begin{aligned} W^+ &\leq \frac{n}{2k}\bar{s} + \frac{1}{v}\sum_{j=1}^{J} p_j E[L_j] \\ &= \frac{n}{2k}\bar{s} + O(\sqrt{n}). \end{aligned} \tag{20}$$

The queue defined by this policy is again a $\sum$GI/G/m queue. Let $\hat{\lambda}_{ij} = \frac{k\lambda}{n}\frac{p_j}{k_j}$ and $\sigma^2_{a_{ij}}$ denote, respectively, the arrival rate and variance of the interarrival time of sets from the i-th subregion of $\mathcal{A}_j$, $i = 1, \ldots, k_j$. Let $\hat{\lambda} = \sum_{j=1}^{J}\sum_{i=1}^{k_j}\hat{\lambda}_{ij} = \frac{k\lambda}{n}$ denote the overall arrival rate of sets. Then, by the same reasoning as in the unbiased case we find that $\sigma^2_{a_{ij}} = \frac{n}{k}(\frac{k_j}{p_j}\sigma_a^2 + \frac{1-p_j/k_j}{(p_j/k_j)^2\lambda^2})$, and therefore

$$\begin{aligned} \sum_{j=1}^{J}\sum_{i=1}^{k_j}\hat{\lambda}^3_{ij}\sigma^2_{a_{ij}} &= \sum_{j=1}^{J}\sum_{i=1}^{k_j}\left(\frac{k\lambda}{n}\frac{p_j}{k_j}\right)^3\left(\frac{n}{k}\right)\left(\frac{k_j}{p_j}\sigma_a^2 + \frac{1-\frac{p_j}{k_j}}{(\frac{p_j}{k_j})^2\lambda^2}\right) \\ &= \hat{\lambda}^2\lambda\sum_{j=1}^{J}\left(p_j(\frac{p_j}{k_j})\sigma_a^2 + \frac{1}{\lambda^2}p_j(1-\frac{p_j}{k_j})\right) \\ &= \hat{\lambda}^2\left(\frac{1}{\lambda} + \frac{\lambda}{k}(\sigma_a^2 - \frac{1}{\lambda^2})\sum_{j=1}^{J}\mu_j^{4/3}A_j\right). \end{aligned}$$

Again, we use the fact that for large values of k, the right hand side above is approximately $\frac{1}{\lambda}\hat{\lambda}^2$. Substituting this approximate expression into Theorem 3 and using Equations (17) and (18) we obtain,

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

Adding the bounds (19) and (20) to the above expression we obtain that as $\rho \longrightarrow 1$,

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

In terms of

$$y = \frac{\lambda\beta\sum_{j=1}^{J}\mu_j^{2/3}A_j}{mv(1-\rho)\sqrt{n}}$$

we have

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

An approximate optimal value for y is

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

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

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

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

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

Comparing this to the lower bound in Theorem 2 establishes the proposition.

□ (Proposition 2)

Again, we remark that a continuous density can be approximated arbitrarily closely by a piecewise uniform density by taking a large number of partitions J above.

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 are provably within a constant factor of the optimal solution. These policies, however, use optimal TSP solutions, which in practice are difficult to compute. In this section, we perform simulation experiments on two other 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 $\mathcal{C} = \{\theta | 0 \leq \theta \leq 1\}$ denote the unit circle and $\mathcal{S} = \{(x, y) | 0 \leq x \leq 1, 0 \leq y \leq 1\}$ denote the unit square. A space filling curve is a continuous mapping ψ from $\mathcal{C}$ onto $\mathcal{S}$ that preserves certain "nearness" properties (*c.f.* Platzman and Bartholdi [14] 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 preimages 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 $\mathcal{C}$.

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

In [5], 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 $\gamma_{SFC} \approx 0.66$ and $\gamma_{NN} \approx 0.64$. In comparison, the modified TSP policy which has $\gamma_{TSP} \approx 0.51$ and the lower bound of Theorem 1 has a value $\gamma_{LB} \approx 0.38$.

We next investigate the *distributional* behavior of these two policies using a similar set of simulation experiments. We show that the SFC policy behaves approximately like a unbiased policy. The NN policy, on the other hand, appears to

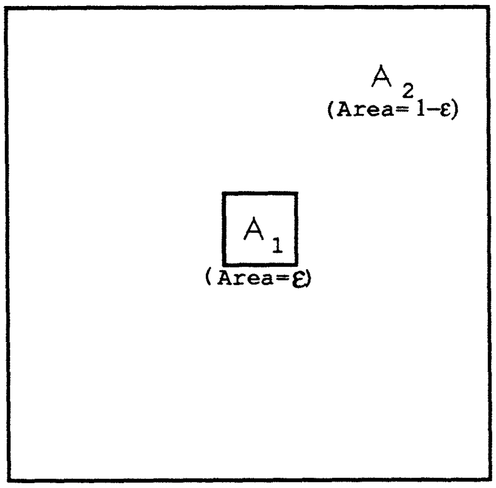

Figure 1: An Extreme Case General Demand Example

behave neither like a pure unbiased policy nor like an optimal biased policies; rather, its performance seems to lie between these two extremes. (We shall quantify these statements below.) 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 $1-\epsilon$ 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

$$f(x) = \begin{cases} \frac{1-\delta}{\epsilon} & x \in \mathcal{A}_1 \\ \frac{\delta}{1-\epsilon} & x \in \mathcal{A}_2 \end{cases} \tag{21}$$

We used identical simulation techniques (*i.e.* same simulation code with different $f(x)$) as in [5]. (See [5] for details.) To estimate the dependence of the system time for each policy, we set $\epsilon = 10^{-4}$ and fixed $A = 1$, $\bar{s} = 0.1$, $\sigma_s^2 = 0$ and $\rho = 0.8$. Then, a different simulation run was performed for eleven 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

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

In the unbiased case $\alpha = 1/2$ and in the optimal biased case $\alpha = 2/3$. For the particular density $f(x)$ given by (21) and for ϵ small,

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

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 $\alpha = 1/2$ for unbiased policies and a value of $\alpha = 2/3$ 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 function of δ for our simulation runs. The 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

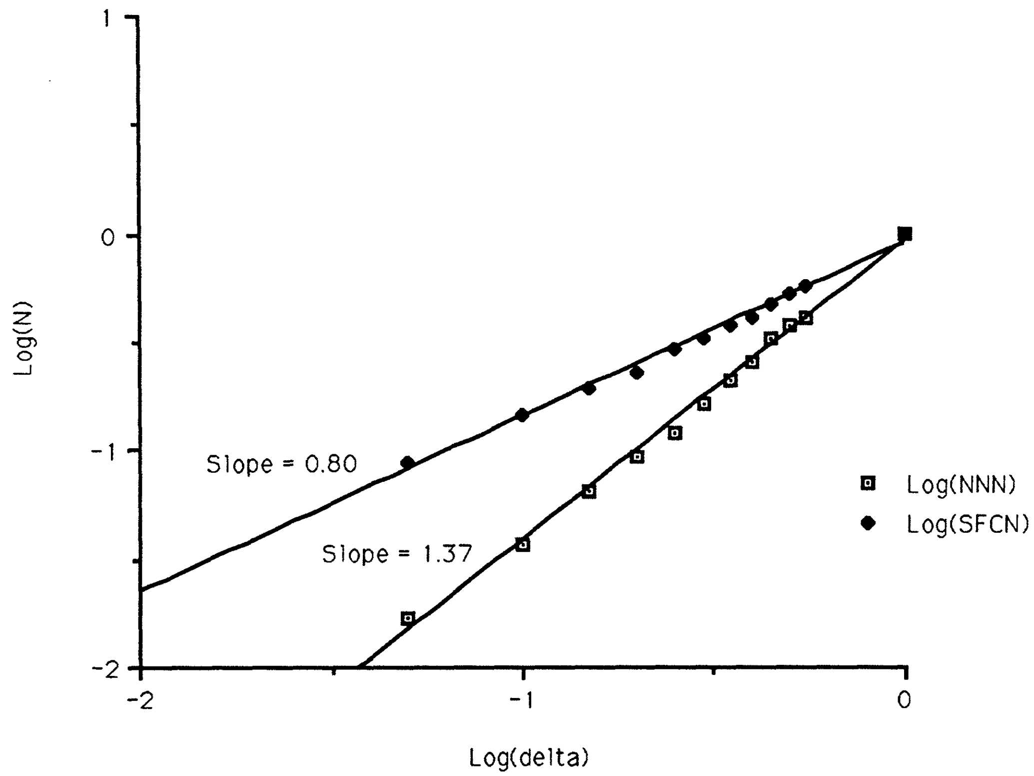

Figure 2: Simulation Results for SFC and NN Policies for General Demand Distribution

to $\alpha = 0.44$ while for the SFC policy, the slope of 1.37 implies $\alpha = 0.58$. These values suggest that the SFC policy performs like a unbiased policy since its value of α is close to 1/2. (Though the performance appears to be somewhat worse (higher α) than a purely unbiased policy.) The NN policy, on the other hand, seems to be between a unbiased and an optimal biased policy; that is, it achieves a higher value of α than a unbiased policy could, but does not achieve as high a value of α as optimally biased policies.

These results also 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

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

where for the uniform demand case, $\Xi = A$, for the spatially unbiased general demand case $\Xi = (\int_{\mathcal{A}} f^{1/2}(x)dx)^2$ and for the spatially biased general demand case $\Xi = (\int_{\mathcal{A}} f^{2/3}(x)dx)^3$. 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 3 *For any continuous density function $f(x)$ defined over the region $\mathcal{A}$ of area A*

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

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 Hardly, Littlewood and Pòlya [8]:

Lemma 4 (Hardly, Littlewood and Pòlya) *If $\alpha > 1$ or $\alpha < 0$, $g(x) \geq 0$ and $h(x) \geq 0$ then*

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

with equality if and only if $\frac{g(x)}{h(x)}$ is constant for all x.

For the first inequality in our proposition, take $g(x) = f(x)$, $h(x) = f^{1/2}(x)$ 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, $(\int_{\mathcal{A}} g(x)dx)^{1-\alpha}(\int_{\mathcal{A}} h(x)dx)^{\alpha} = (\int_{\mathcal{A}} f^{1/2}(x)dx)^2$. Thus,

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

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

For the second inequality, take $g(x) = f^{2/3}(x)$, $h(x) = f^{1/2}(x)$ 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

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

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

□ (Proposition 3)

Proposition 3 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 nonuniform distribution f. Also, note that when the density is uniform there is nothing to be gained by not providing unbiased 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 a 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 of 1. For the density used in this example, it is straightforward to show that for a fixed $\delta > 0$ and $\epsilon \longrightarrow 0$,

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

and

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

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

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

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 a 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 but 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 $\mathcal{A}_2$ will then build up and thus can be serviced more efficiently with larger tours. This frees up more vehicle time to service the much higher fraction of customers 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 2, one can see that very little of the vehicle routing "structure" inherent in the DTRP was used. Indeed, we only assumed that the service was sequential (*i.e.* one demand served at a time), which allowed us the establish that the mean number left behind by a departure from any given region was 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 2. 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 2 is in fact tight within this broader class of *removal rules*, and therefore more vehicle routing features of the DTRP need to be exploited in 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.

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 $\frac{\pi z^2}{A}$ (ignoring edge effects because z is small) and the mean interarrival time of points is $\frac{1}{\lambda}$, 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 points location is uniformly distributed within the circle of radius z about x_i, we have

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

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 2 and recalling that $\phi(x) = f(x) = 1/A$ for the spatially unbiased, uniform case shows that the lower bound is indeed tight within the class of sequential removal rules.

The result can be extended to the nonuniform case by taking the radius z above to be a function of a points location x; that is, $z(x)$. Define

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

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

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

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,

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

Unconditioning implies

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

which establishes the tightness of the lower bound for the general unbiased case as well. (We do not have an analogous example for the general biased case.)

6.2 Relation to the Static Nearest Neighbor Bound

The bound in Lemma 2 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 such Euclidean problems such as the TSP, Matching and Minimum Spanning Tree [13]. In the same sense that this nearest neighbor bound is weak for the static TSP, one can see that the bound of Lemma 2 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 too pessimistic. Indeed, we *conjecture* that these policies are in fact asymptotically optimal.

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 form 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 [5] and [6].

By simply using the more general bound on the nearest neighbor distance $E[Z^*]$ of Lemma 2 in the arguments of [6], one can show the following theorems:

Theorem 4 *Within the class of spatially unbiased policies*

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

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

Theorem 5 *Within the class of spatially biased policies*

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

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

A provably good unbiased policy for the finite capacity case can be obtained by modifying the unbiased policy from Proposition 1 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 [7] as was done for the Modified qTP policy in [6]. Serve these sets FCFS and optimize over n. For large k, the resulting system time, T_{qU}, then satisfies

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

which implies the same performance guarantee as in the uniform case (*c.f.* [6]).

An identical tour partitioning modification applied to the sets formed in the spatially biased policy of Proposition 2 gives a policy with a system time, T_{qB}, satisfying

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

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 Higher Dimensions

Most of the DTRP bounds and policies can be extended to Euclidean subsets $\mathcal{A}$ of $\Re^d$ 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 6 *Within the class of spatially unbiased policies*

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

where $\gamma(d) = \frac{d}{d+1}\left(\frac{1}{d+1}\right)^{1/d}\left(\frac{1}{c_d}\right)^{1/d}$ *and* $c_d = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2}+1)}$ *is the volume of the ball of unit radius in dimension* d *(*$\Gamma(x)$ *is the usual gamma function).*

Theorem 7 *Within the class of spatially biased policies*

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

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 $(1-\rho-\frac{2\lambda\overline{r}}{mvq})$ 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

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

where $\Xi(d) = V$ for the uniform case, $\Xi(d) = \left[\int_{\mathcal{A}} f^{\frac{d-1}{d}}(x)dx\right]^d$ for the spatially unbiased case and $\Xi(d) = \left[\int_{\mathcal{A}} f^{\frac{d}{d+1}}(x)dx\right]^{d+1}$ for the spatially biased case. For example, the modified TSP policy in d dimensions has a constant value of $\frac{\beta(d)}{2^{1/d}}$, where $\beta(d)$ is the d dimensional TSP constant.

An interesting result is found by examining this policy for $d\to\infty$. In [4], it was conjectured and subsequently proved in [16] that for $d\to\infty$

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

By using the fact that for $d\to\infty$, $\frac{d}{d+1}\sim 1$, $\left(\frac{1}{d+1}\right)^{1/d}\sim 1$ and $\Gamma(\frac{d}{2}+1)\sim\sqrt{2\pi}(\frac{d}{2})^{\frac{d}{2}+\frac{1}{2}}e^{-\frac{d}{2}}$, it is straightforward to show that

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

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

Theorem 8 *For the uncapacitated, m-server DTRP, the modified TSP policy is an optimal heavy traffic policy asymptotically as* $d\to\infty$.

This theorem gives further evidence of the asymptotic optimality of the modified TSP policy, which we have already 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 customer locations and biasness 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 reoccurring 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 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 can form the basis for effective policies in dynamic, stochastic environments.

References

[1] Bartholdi, J.J. and Platzman, L.K. (1982), "An O(NlogN) Planar Traveling Salesman Heuristic Based on Space Filling Curves," *Operations Research Letters*, 1, 4, 121-125.

[2] Bartholdi, J.J. and Platzman, L.K. (1988), "Heuristics Based on Spacefilling Curves for Combinatorial Problems in Euclidean Space", *Mgmt. 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 van Ryzin, G. (1990), "An Asymptotic Determination of the Minimum Spanning Tree and Minimum Matching Constants in Geometrical Probability," *Operations Research Letters*, 9, 223-231.

[5] Bertsimas, D. and van Ryzin, G. (1990), "A Stochastic and Dynamic Vehicle Routing Problem in the Euclidean Plane", to appear in *Operations Research.*

[6] Bertsimas, D. and van Ryzin, G. (1991), "Stochastic and Dynamic Vehicle Routing in the Euclidean Plane with Multiple Capacitated Vehicles", submitted to *Operations Research.*

[7] Haimovich, M and Rinnoy Kan, A. H. G. (1985), "Bounds and Heuristics for Capacitated Routing Problems," *Mathematics of Operations Research*, 10, 4, 527-542.

[8] Hardy, G. H., Littlewood, J. E. and Pòlya, G. (1934), *Inequalities*, Cambridge.

[9] Inglehart, D. L. and Whitt, W. (1970), "Multiple Channel Queues in Heavy Traffic, I and II," *Advances in Applied Probability*, 2, 150-177 and 355-364.

[10] Johnson, D. (1988), talk presented at the Mathematical Programming Symposium, Tokyo.

[11] Karp R., (1977), "Probabilistic Analysis of Partitioning Algorithms for the Traveling Salesman in the Plane", *Math. Oper. Res.*, 2, 209-224.

[9] Inglehart, D. L. and Whitt, W. (1970), "Multiple Channel Queues in Heavy Traffic, I and II," *Advances in Applied Probability*, 2, 150-177 and 355-364.

[10] Johnson, D. (1988), talk presented at the Mathematical Programming Symposium, Tokyo.

[11] Karp R., (1977), "Probabilistic Analysis of Partitioning Algorithms for the Traveling Salesman in the Plane", *Math. Oper. Res.*, 2, 209-224.

[12] Lawler, E.L., J.K. Lenstra, A.H.G. Rinnooy Kan, D.B. Shmoys (eds) (1985), *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*, Wiley, Chichester.

[13] Papadimitriou, C. H., "The Probabilistic Analysis of Matching Heuristics," in Proc. Fifteenth Annual Allerton Conf. on Communication, Control and Computing, 368-378.

[14] L.K. Platzman and J.J. Bartholdi (1989), "Spacefilling Curves and the Planar Travelling Salesman Problem", *JACM*), 719-737.

[15] Psaraftis, H. (1988) "Dynamic Vehicle Routing Problems", in *Vehicle Routing: Methods and Studies* (B. Golden, A. Assad, eds.), North Holland.

[16] Rhee, W., "On the TSP in Many Dimensions," Unpublished Manuscript, 1991.

[17] Wolff, R. W. (1982), "Poisson Arrivals See Time Averages," *Operations Research*, 30, 2, 223-231.