**A Stochastic and Dynamic Vehicle Routing Problem
in the Euclidean Plane**

Dimitris J. Bertsimas
Garrett van Ryzin

A Stochastic and Dynamic Vehicle Routing Problem in the Euclidean Plane

Dimitris J. Bertsimas *   Garrett van Ryzin [tt]

February 1990

Abstract

We propose and analyze a generic mathematical model for dynamic, stochastic vehicle routing problems which we call the dynamic traveling repairman problem (DTRP). The model is motivated by applications in which the objective is minimizing the wait for service in a stochastic and dynamically changing environment. This is a departure from traditional vehicle routing problems which seek to minimize total travel time in a static, deterministic environment. Potential areas of application include repair, inventory, emergency service and scheduling models. The DTRP is defined as follows: Demands for service arrive according to a Poisson process in a Euclidean region $\mathcal{A}$ and, upon arrival, are independently and uniformly assigned a location in $\mathcal{A}$. Each demand requires an independent and identically distributed service by a vehicle that travels at unit velocity. The problem is to find a policy that minimizes the average time a demand spends in the system. We propose several policies for the DTRP and analyze their behavior. Using approaches from queueing theory, geometrical probability, combinatorial optimization and simulation, we find a provably optimal policy in light traffic and several policies that have system times within a constant factor of the optimal policy in heavy traffic. We also show that the waiting time grows much faster than in traditional queues as the traffic intensity increases, yet the stability condition does not depend on the system geometry.

1

Introduction

The traveling salesman problem (TSP) is one of the most studied problems in the operations research and applied mathematics literature. The attention it receives is due both to the problem's richness and inherent elegance and to the frequent occurrence of the TSP in practical problems, both directly and as a subproblem. Yet, in many practical applications, the TSP is a deterministic, static approximation to a problem which is in reality both probabilistic and time varying (dynamic). In addition there are often costs associated with the wait for delivery that are not captured in the objective of minimizing travel distance.

For example, a prototypical application of the TSP is the routing of a vehicle from a central depot to a set of dispersed demand points so as to minimize the total travel costs. In real distribution systems, however, orders (demands) arrive randomly in time, and the dispatching of vehicles is a continuous process of collecting demands, forming tours and dispatching vehicles. In such a dynamic setting, the wait for a delivery (service) may be a *more* important factor than the travel cost. Applications in which the wait for service rather than the total travel time is a more suitable objective and also the demand pattern is both dynamic and stochastic include the following:

1. The demands are requests for replenishment of stock (raw materials, merchandise, etc.) from remote sites that must be delivered from a central depot. In this case, large waiting times mean that large inventories are needed at the remote sites to prevent stock-out.

2. In managing a fleet of taxis, one would like to minimize the average waiting time of customers. Decision makers therefore need good dispatching policies, fleet sizing models and estimates of the level of service.

3. Demands represent requests for emergency service. The objective is therefore to reduce the wait for service rather than to minimize the travel cost of the emergency vehicle. In this case, we want real-time policies that can be applied in a stochastic environment.

4. The demands are geographically dispersed failures that must be serviced by a mobile repairman. The objective in this case is to minimize the downtime (wait plus service time) at the various locations. Examples in this category include servicing of geographically distributed communications or utility networks, automobile road service

3

(AAA), or the dispatching of a roving expert to local sites.

5. Finally for completeness, consider the problem in which a salesman receives leads randomly in time and wants to make his sales calls so as to minimize the time customers spend contemplating their purchases!

Motivated by these application areas, we propose and analyze a generic mathematical model which we call the dynamic traveling repairman problem (DTRP). The model has several important characteristics:

1. The objective is to minimize waiting time not travel cost.

2. Information about future demands is stochastic.

3. The demands vary over time (i.e. they are dynamic).

4. Policies have to implemented in real time.

5. The problem involves queueing phenomena.

In general, little is know about dynamic versions of the TSP. Psaraftis [25] defines the dynamic traveling salesman problem (DTSP), which initially motivated our investigation: In a complete graph on n nodes, demands for service are independently generated at each node i according to a Poisson process with parameter λ_i. These demands are to be serviced by a salesman who takes a known time t_{ij} to go from i to j, and spends a stochastic time X, which has a known distribution, servicing each demand (on location). The goal is to find strategies that optimize over some performance measure (waiting time, throughput). By comparison, the DTRP is defined in the Euclidean plane and optimization is over the total system time. No general results were obtained for this problem, but useful insights and conjectures were made.

Some of the above characteristics of the DTRP have also been considered in isolation in the literature. The first is the objective of minimizing average system time rather than total travel time. In a deterministic setting, this idea appears in the traveling repairman (or delivery) problem (TRP), in which a repair unit has to service a set of demands V starting from a depot. If $d(i, j)$ denotes the travel time from i to j, the problem is to find a tour starting from the depot through the demands so as to minimize the total waiting time of the demands. As a result, if the sequence in which the repair unit travels is $t = (1, 2, \ldots, n, 1)$

then the total waiting time is $W_t = \sum_{i=1}^n w_i$, where $w_i = \sum_{j=1}^{i-1} d(j, j+1)$ is the waiting time of the demand i. The problem closely resembles the TSP and can be thought of as the deterministic and static analog of the DTRP. As is the case with the TSP, the TRP is NP-complete both on a graph and in the Euclidean plane (Sahni and Gonzalez [27]). In contrast with the TSP, which is trivial on trees, the TRP seems difficult on trees. Minieka [23] proposes an exponential $O(n^p)$ algorithm for the TRP on a tree $T = (V, E)$, where $|V| = n$ and p is the number of leaves in T. Despite its interest and applicability the problem has not received much attention from the research community. As a result not much is known about the TRP.

Jaillet [13], Bertsimas [7] and Bertsimas, Jaillet and Odoni [8] address the second and fourth characteristics under the unifying framework of a priori optimization. They define and analyze the probabilistic traveling salesman problem (PTSP) and the probabilistic vehicle routing problem (PVRP) as follows: There are n known points, and on any given instance of the problem only a subset S consisting of $|S| = k$ out of n points $(0 \leq k \leq n)$ must be visited. Suppose that the probability that instance S occurs is $p(S)$. We wish to find a priori a tour through all n points. On any given instance of the problem, the k points present will then be visited in the same order as they appear in the a priori tour. The problem of finding such an a priori tour which is of minimum length in the expected value sense is defined as the PTSP. In the case were the vehicle has capacity Q, the corresponding problem is the probabilistic vehicle routing problem. It is clear that the policy followed is a real-time policy, but the problem is inherently static and is solved a priori using only probabilistic information.

An important characteristic of the DTRP is that it incorporates queueing phenomena. Queueing considerations in the context of location problems have been considered in Berman et. al. [4] and Batta et. al. [2]. In this setting, the authors define the stochastic queue median problem (SQMP) in which the important decision is a strategic one; we would like to locate a server in a network which behaves like an $M/G/1$ queue. Arrivals occur in a dynamic manner according to a Poisson process, and a server (vehicle), following a first-come-first-serve (FCFS) discipline, is dispatched from a central depot and then returns to the depot again after service is completed. The problem is to locate the depot on a network so that the mean queueing delay plus travel time is minimized. The model is very appropriate for analyzing emergency service systems (e.g. police, fire and ambulance

5

service). In our setting, the SQMP can be seen as a network case of the DTRP in which the policy followed is to strategically locate the server and then follow a FCFS dispatching rule. The connection between location/queueing problems and dynamic vehicle routing problems was also recognized by Psaraftis [25] where he conjectures that for low arrival rates, the DTSP resembles the 1-median problem. We analyze the performance of this policy in Section 4.1 for the Euclidean case and show formally that this is indeed true.

Our strategy in analyzing the DTRP is the following: First, we establish some lower bounds on the average system time for all policies. Then, using a variety of techniques from combinatorial optimization, queueing theory, geometrical probability and simulation, we analyze several policies and compare their performance to the lower bounds. A variant of the FCFS policy, called the stochastic queue median policy, is shown to be optimal in the case of light traffic. In heavy traffic, several policies are shown to be within a constant factor of the lower bounds and thus from the optimal policy. The policy with the best provable performance guarantee in heavy traffic is one based on forming TSP tours, while the best policy empirically is the nearest neighbor policy. Our results also show that the system time grows much more rapidly with traffic intensity than in traditional queues and that the stability condition is independent of the system geometry.

The paper is organized as follows: Since we use a variety of results from several areas, we briefly describe them and give appropriate references in Section 1. In Section 2, we formally describe the DTRP and introduce notation. Lower bounds for the optimal system time are derived in Section 3. In Section 4, which is central to the paper, we introduce and analyze several policies for the DTRP. In Section 4.6, an example is given to illustrate the relative performance of the policies. Finally in Section 5 we summarize the contributions of the paper and give some concluding remarks.

1 Probabilistic and Queueing Background

In this section, we briefly describe the results used in the following sections of the paper.

An Upper Bound for the Waiting Time in a $GI/G/1$ Queue

In a $GI/G/1$ queue let $\frac{1}{\lambda}, \bar{s}$ be the expected interarrival and service time and let σ_a^2, σ_s^2 the variances of the interarrival and service time distribution respectively. Let $\rho = \lambda \bar{s}$ be the traffic intensity. There is no simple explicit expression for the expected waiting time W in

this case. (The average system time T is simply $W + \bar{s}$.) Kingman [16] (see also Kleinrock [18]) proves that

$$W \leq \frac{\lambda(\sigma_a^2 + \sigma_s^2)}{2(1-\rho)}. \tag{1}$$

In addition, this upper bound is asymptotically exact as $\rho \to 1$. For the $M/G/1$ it is well known (see Kleinrock [18]) that

$$W = \frac{\lambda \overline{s^2}}{2(1-\rho)}, \tag{2}$$

where $\overline{s^2} = \sigma_s^2 + \bar{s}^2$ is the second moment of the service time.

Symmetric Cyclic Queues

Consider a queueing system that consists of k queues $Q_1, Q_2, \ldots, Q_k$ each one with infinite capacity. Customers arrive at each queue according to independent Poisson processes with the same arrival intensity λ/k. The queues are served by a single server who visits the queues in a fixed cyclic order $Q_1, Q_2, \ldots, Q_k, Q_1, Q_2, \ldots$. The travel time around the cycle is d. The service times at every queue are independent, identically distributed random variables with mean $\bar{s}$ and second moment $\overline{s^2}$. The traffic intensity is $\rho = \lambda \bar{s}$. The server uses the exhaustive service policy, i.e. servicing each queue i until the queue is empty before proceeding. The expected waiting time for this system is given by (see Bertsekas and Gallager [5], p.156)

$$W = \frac{\lambda \overline{s^2}}{2(1-\rho)} + \frac{(1 - \frac{\rho}{k})}{2(1-\rho)} d. \tag{3}$$

We note that in an asymmetric cyclic queue, in which arrival processes and service times are not identical, there are no closed form expressions for the waiting time (see Ferguson and Aminetzah [11]).

Jensen's Inequality

If f is a convex function and X is a random variable then

$$E[f(X)] \geq f(E[X]), \tag{4}$$

provided the expectations exist.

Markov's Inequality

If X is a nonnegative random variable and x is any nonnegative number, then

$$E[X] \geq xP\{X > x\}. \tag{5}$$

Wald's Equation

Let $\{X_i; i \geq 1\}$ be a sequence of i.i.d. random variables with $E[X] < \infty$ and N be a finite

mean random variable with the property that $P\{N = n\}$ is independent of $\{X_i; i > n\}$ for all n. (Such a random variable N is said to be a *stopping time* for the sequence $\{X_i; i \geq 1\}$.) Then

$$E\left[\sum_{i=1}^{N} X_i\right] = E[N]E[X]. \tag{6}$$

Stochastically Larger (Definition)

A random variable X is said to be *stochastically larger* than a random variable Y, denoted $X \geq_{ST} Y$, if

$$P\{X > z\} \geq P\{Y > z\} \quad \text{for all } z. \tag{7}$$

Geometrical Probability

Given two uniformly and independently distributed points X_1, X_2 in a square of area A, then

$$E[\|X_1 - X_2\|] = c_1 \sqrt{A}, \quad E[\|X_1 - X_2\|^2] = c_2 A, \tag{8}$$

where $c_1 \approx 0.52, c_2 = \frac{1}{3}$ (see Larson and Odoni [19], p.135). If we let x^* denote the center of a square of area A, then it is known [19] that the first and second moment of the distance to a uniformly chosen point X are given by

$$E[|X - x^*|] = c_3 \sqrt{A}, \quad E[|X - x^*|^2] = c_4 A, \tag{9}$$

where $c_3 = (\sqrt{2} + ln(1 + \sqrt{2})/6 \approx 0.383, c_4 = \frac{1}{6}$.

Asymptotic Properties of the TSP in the Euclidean Plain

Let $X_1 \ldots X_n$ be independently and uniformly distributed points in a square of area A and let L_n denote the length of the optimal tour through the points. Then there exists a constant β_{TSP} such that

$$\lim_{n \to \infty} \frac{L_n}{\sqrt{n}} = \beta_{TSP} \sqrt{A}. \tag{10}$$

with probability one (see [28], [22]). In his recent experimental work with very large scale TSP's, Johnson [14] estimated $\beta_{TSP} \approx 0.72$. In addition, it is also well known (see [22], p. 189) that $\lim_{n \to \infty} var(L_n) = O(1)$, and therefore

$$\lim_{n \to \infty} \frac{var(L_n)}{n} = 0. \tag{11}$$

Space Filling Curves

The following results are due to Platzman and Bartholdi [24]. 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. Then there exists a continuous mapping ψ from $\mathcal{C}$ onto $\mathcal{S}$ with the property that for any $\theta, \theta' \in \mathcal{C}$,

$$\|\psi(\theta) - \psi(\theta')\| \leq 2\sqrt{|\theta - \theta'|}. \tag{12}$$

If $X_1 \ldots X_n$ are any n points in $\mathcal{S}$ and L_n is the length of a tour of these n points formed by visiting them in increasing order of their *preimages* in $\mathcal{C}$ (i.e. increasing θ order), then

$$L_n \leq 2\sqrt{n}. \tag{13}$$

If the points $X_1 \ldots X_n$ are independently and uniformly distributed $\mathcal{S}$, then there exists a constant β_{SPC} such that

$$\limsup_{n \to \infty} \frac{L_n}{\sqrt{n}} = \beta_{SFC} \tag{14}$$

with probability one. The value of β_{SFC} is approximately 0.956.

2 Problem Definition and Notation

The DTRP is defined as follows: A convex region $\mathcal{A}$ of area A contains a vehicle (server) that travels at constant, unit velocity between demand (or customer) locations. To simplify the calculations and the presentation, in most cases we assume the region $\mathcal{A}$ is a square of area A. This restriction can often be relaxed without affecting the results, though actual calculations may be more difficult. Demands for service arrive according to a Poisson process with rate λ and upon arrival are independently and uniformly assigned a location within $\mathcal{A}$. Each demand i requires an independent and identically distributed on-site service with mean duration $\bar{s}$ and second moment $\overline{s^2}$. It is assumed that $\bar{s} > 0$. The fraction of time the server spends in on-site service is denoted ρ, and for stable systems $\rho = \lambda\bar{s}$.

The travel time *to* the ith demand location, d_i, is defined as follows: if the ith demand is the first to be serviced in a busy period, then d_i is the travel time from the server's location at the time of demand $i's$ arrival to demand $i's$ location; otherwise, d_i is the travel time from the location of the demand served prior to i to demand $i's$ location. The term d_i can be considered the travel time component of demand $i's$ *total service requirement*. The steady state expected value of d_i is denoted $\bar{d}$ and is given by $\bar{d} \equiv \lim_{i \to \infty} E[d_i]$, where we assume the limit exists.

The system time of demand i, denoted T_i, is defined as the elapsed time between the arrival of demand i and the time the server completes the service of i. The waiting time

of demand i, W_i is defined by $W_i = T_i - s_i$. The steady state system time, T, is define by $T \equiv \lim_{i \to \infty} E[T_i]$ and $W \equiv T - \bar{s}$. The problem is to find a policy for servicing demands that minimizes T, and this optimal system time is denoted T^*. We use system time rather than waiting time because in relating the DTRP to traditional queueing systems d_i can be mistakenly interpreted as part of the "service time", which does not correspond to the definition above.

A final remark concerning the difference between the DTRP and the $M/G/1$ queue. In the DTRP, the total service requirement has both a travel and on-site service component. Although the on-site service requirements are independent, the travel times are generally not. As a result, total service requirements are not i.i.d. random variables, and therefore the methodology of the $M/G/1$ queue is not directly applicable.

3 Lower Bounds on the Optimal DTRP Policy

We first establish two simple but powerful lower bounds on the optimal expected system time, T^*. In Section 4, we then use these lower bounds to evaluate the performance of the proposed policies.

3.1 A Light Traffic Lower Bound

The first bound for the DTRP is established by dividing the system time of customer i, T_i, into three components: the waiting time of customer i due to the servers travel prior to service of i, denoted W_i^d; the waiting time of customer i due to service of customers served prior to customer i, denoted W_i^s; and customer $i's$ service time, s_i. Thus,

$$T_i = W_i^d + W_i^s + s_i.$$

Taking expectations and letting $i \to \infty$ gives

$$T = W^d + W^s + \bar{s}, \tag{15}$$

where $W^d \equiv \lim_{i \to \infty} E[W_i^d]$ and $W^s \equiv \lim_{i \to \infty} E[W_i^s]$. Note that $W = W^d + W^s$.

To bound W^d, note that the travel component of the waiting time of a given customer (demand) is at least the travel delay between the servers location at the time of the customer's arrival and the customer's location. In general, the server is located in the region

10

according to some (generally unknown) spatial distribution that depends on the server's policy; thus, W^d is bounded by the expected distance from a server location selected from this distribution to a uniform location. Now, suppose we had the option of locating the server in the best *a priori* location, x^*; that is, the location that minimizes the expected distance to a uniformly chosen location, X. This certainly yields a lower bound on the expected distance between the server and the arrival, so

$$W^d \geq \min_{x_0 \in \mathcal{A}} E[\|X - x_0\|]. \tag{16}$$

The location x^* that achieves the minimization above is the *median* of the region $\mathcal{A}$. For the case where $\mathcal{A}$ is a square, x^* is simply the center of the square, in which the lower bound is from (9),

$$W^d \geq c_3 \sqrt{A} \approx 0.383 \sqrt{A}. \tag{17}$$

To bound W^s, let N denote the expected number of customers served during a customer's waiting time. Then since service times are independent, we have

$$W^s = \bar{s} N + \frac{\lambda \overline{s^2}}{2},$$

where the second term is the expected residual service time of the customer in service. Since in steady state the expected number of customers served during a wait is equal to the expected number who arrive, we can apply Little's law to get

$$W^s = \bar{s} \lambda W + \frac{\lambda \overline{s^2}}{2} = \rho W + \frac{\lambda \overline{s^2}}{2}.$$

Since $W = W^d + W^s$ we obtain

$$W^s = \frac{\rho}{1 - \rho}(W^d) + \frac{\lambda \overline{s^2}}{2(1 - \rho)}. \tag{18}$$

Combining (15), (16) and (18) and noting that these bounds are true for *all* policies we get the following theorem:

Theorem 1

$$T^* \geq \frac{E[\|X - x^*\|]}{1 - \rho} + \frac{\lambda \overline{s^2}}{2(1 - \rho)} + \bar{s}. \tag{19}$$

where x^* is the median of the region $\mathcal{A}$. For the special case where $\mathcal{A}$ is a square,

$$E[\|X - x^*\|] = c_3 \sqrt{A} \approx 0.383 \sqrt{A}.$$

As shown below, this bound is most useful in the case of light traffic ($\lambda \to 0$).

11

3.2 A Heavy Traffic Lower Bound

A lower bound that is most useful for $\rho \to 1$ is provided by the following theorem:

Theorem 2 *There exists a constant γ such that*

$$T^* \geq \gamma^2 \frac{\lambda A}{(1-\rho)^2} - \frac{1-2\rho}{2\lambda}. \tag{20}$$

Proof

First, suppose that for *all* service policies the following bound is known

$$\bar{d} \geq \gamma \frac{\sqrt{A}}{\sqrt{N+1/2}}, \tag{21}$$

where N is the average number of customers in queue. Then the stability condition,

$$\bar{s} + \bar{d} \leq \frac{1}{\lambda}, \tag{22}$$

implies

$$\bar{s} + \frac{\gamma\sqrt{A}}{\sqrt{N+1/2}} \leq \frac{1}{\lambda}.$$

After rearranging, noting that $T = W + \bar{s}$ and $N = \lambda W$, we obtain the bound of Theorem 2. So the theorem is established once (21) is proven.

To prove (21), consider a random "tagged" arrival and define,

$\mathcal{S}_0$: *The set of locations of customers who are in queue at the time of the tagged customer's arrival union with the server's location.*

$\mathcal{S}_1$: *The set of locations of the customers who arrive during the tagged customer's waiting time ordered by their time of arrival.*

$X_0 \equiv$ *The tagged customers location.*

$N_i \equiv |\mathcal{S}_i| \quad i = 0, 1$

$Z_0^* \equiv \min_{x \in \mathcal{S}_0} \|x - X_0\|$

Further, define $Z_i \equiv \|X_i - X_0\|$ where X_i is the location of the ith customer to arrive after the tagged customer (e.g. $\mathcal{S}_1 = \{X_1, X_2, ..., X_{N_1}\}$). Note that $\{Z_i; i \geq 1\}$ are i.i.d. with

$$P\{Z_i \leq z\} \leq \frac{\pi z^2}{A}, \tag{23}$$

12

and that N_1 is a stopping time for the sequence $\{Z_i; i \geq 1\}$.

The set of locations from which the server can visit the tagged customer is at most $\mathcal{S}_0 \cup \mathcal{S}_1$; therefore, the travel time component of the tagged customer's total service requirement is at least $\min\{Z_0^*, Z_1, ..., Z_{N_1}\}$. Hence,

$$\bar{d} \geq E[\min\{Z_0^*, Z_1, ..., Z_{N_1}\}]. \tag{24}$$

We next bound the right hand side of (24). To do so define a indicator variable for a random variable X by

$$I_X = \begin{cases} 1 & \text{if } X \leq z \\ 0 & \text{if } X > z \end{cases}$$

where z is a positive constant to be determined below. Then

$$
\begin{aligned}
P\{\min\{Z_0^*, Z_1, ..., Z_{N_1}\} > z\} &= P\{I_{Z_0^*} + \sum_{i=1}^{N_1} I_{Z_i} = 0\} \\
&= 1 - P\{I_{Z_0^*} + \sum_{i=1}^{N_1} I_{Z_i} > 0\} \\
&\geq 1 - E[I_{Z_0^*} + \sum_{i=1}^{N_1} I_{Z_i}] \qquad (I_X \ Integer) \\
&= 1 - E[I_{Z_0^*}] - E[N_1]E[I_{Z_i}] \qquad (Wald's \ Eq.).
\end{aligned}
$$

Using Markov's inequality and the fact that $E[I_X] = P\{X \leq z\}$, yields

$$E[\min\{Z_0^*, Z_1, ..., Z_{N_1}\}] \geq z(1 - P\{Z_0^* \leq z\} - E[N_1]P\{Z_i \leq z\}).$$

Since $E[N_1] = N$ and $P\{Z_i \leq z\}$ is bounded according to (23), we obtain

$$E[\min\{Z_0^*, Z_1, ..., Z_{N_1}\}] \geq z(1 - P\{Z_0^* \leq z\} - N\frac{\pi z^2}{A}). \tag{25}$$

An upper bound on $P\{Z_0^* \leq z\}$ is provided by the following lemma.

Lemma 1 : $P\{Z_0^* \leq z\} \leq \frac{\pi z^2}{A}(N + 1)$.

Proof

The proof relies on the following result due to Haimovich and Magnanti [12] for the k-median problem: Let $\mathcal{S}$ be *any* set of points in $\mathcal{A}$ with $|\mathcal{S}| = k$, X be a uniformly distributed location in $\mathcal{A}$ independent of $\mathcal{S}$ and define $Z^* \equiv \min_{x \in \mathcal{S}} \|X - x\|$. Define the random variable

13

Y to be the distance from the center of a circle of area A/k to a uniformly distributed point within the circle. Then for all nondecreasing functions f,

$$E[f(Z^*)] \geq E[f(Y)].$$

An immediate consequence of this (see, for example [26]) is that $Z^* \geq_{ST} Y$. As a result,

$$P\{Z^* \leq z\} \leq P\{Y \leq z\} \leq \frac{\pi z^2}{A} k,$$

where the last inequality follows from the definition of Y.

Now consider conditioning on N_0, and note that X_0 is independent of $\mathcal{S}_0$ under any condition on $\mathcal{S}_0$. Therefore, from the above result

$$P\{Z_0^* \leq z | N_0\} \leq \frac{\pi z^2}{A} N_0.$$

Unconditioning and observing that $E[N_0] = N + 1$ establishes the lemma.

$$\square (Lemma\ 1)$$

Using the result of lemma in (25) yields

$$E[\min\{Z_0^*, Z_1, ..., Z_{N_1}\}] \geq z(1 - \frac{\pi}{A}(2N + 1)z^2).$$

Maximizing the right hand side with respect to z gives

$$E[\min\{Z_0^*, Z_1, ..., Z_{N_1}\}] \geq \frac{2}{3\sqrt{6\pi}} \frac{\sqrt{A}}{\sqrt{N + 1/2}}.$$

This establishes (21) with $\gamma = \frac{2}{3\sqrt{6\pi}} \approx 0.153$, and thus the theorem is proved.

$$\square (Theorem\ 2)$$

A few comments on the lower bound of Theorem 2 are in order. First, it shows that the waiting time grows at least as fast as $(1 - \rho)^{-2}$ rather than $(1 - \rho)^{-1}$ as is the case for a classical queueing system. Also, it is only a function of the first moment of the on-site service time, which is again a significant departure from traditional queueing system behavior (e.g. the $M/G/1$ system). The explanation lies in the geometry of the system. The bound of Theorem 2 gives (via Little's Theorem) the minimum average number of customers that must be maintained in the system to ensure that the average travel distance, $\overline{d}$, satisfies the stability condition (22). This number, however, grows much more rapidly than the average number in the system due simply to traditional queueing delays.

Because several optimistic assumptions were used in the above proof (e.g. $\mathcal{S}_0$ is the set of N_0-median locations), it seems that the value $\gamma \approx 0.153$ is probably excessively

low. For example, if one assumes locations of customers at service completion epochs are approximately uniform, then by a modified argument a value of $\gamma = 1/2$ is obtained. We conjecture that Theorem 2 remains true for this larger value of γ.

4   Some Proposed Policies for the DTRP

In this section, we propose and analyze several policies for the DTRP. The first class of policies are based on variants of the FCFS discipline. We show that one such policy is optimal in light traffic, in the sense that it asymptotically achieves the light traffic lower bound of the last section for $\lambda \rightarrow 0$. These policies, however, are unstable for high utilizations; therefore, we turn next to a partitioning policy based on subdividing the large square $\mathcal{A}$ into smaller squares, each of which is served locally using a FCFS discipline. Using results on cyclic queues, we show that the this policy is within a constant factor of the lower bounds for all values of $\rho < 1$. This also establishes $\rho < 1$ as the stability condition in the sense that there exists stable policies for every $\rho < 1$. We next introduce a more sophisticated policy based on forming successive TSP tours. Its average system time is nearly half that of the partitioning policy. Next, a policy based on space filling curves is examined. It too has a constant factor performance guarantee and is shown via simulation to perform about 15% better than the TSP policy. Finally, we examine the policy of serving the nearest neighbor. Because of analytical difficulties, we simulate it and show the average system time is about 10% lower than the SFC policy.

4.1   FCFS Policies

The simplest policy for the DTRP is to service customers in the order in which they arrive (FCFS). The first policy we examine of this type is defined as follows: 1) when customers are present, the server travels directly from one customer location to the next following a FCFS order, and 2) when no customers are present at a service completion, the server waits until the next customer arrives before moving. The random variable d_i is, therefore, the distance between two independent, uniformly distributed locations in $\mathcal{A}$.

Because customer locations are independent of the order of arrivals and also the number of customers in queue, the system behaves like an $M/G/1$ queue. Note that the travel times d_i are not strictly independent (e.g. consider the case $d_i = \sqrt{2A}$); however, it is true that

15

they are identically distributed and also independent of the number in queue. Therefore, the Pollaczek-Khinchin (P-K) formula (2) still holds. (See [5] page 142-143 for a proof of the P-K formula that does not require mutual independence of service times.)

The first and second moments of the total service requirement are, by (8), $\bar{s} + c_1\sqrt{A}$ and $\overline{s^2} + 2c_1\sqrt{A}\bar{s} + c_2A$ respectively, where $c_1 \approx 0.52, c_2 = \frac{1}{3}$. The average system time is therefore

$$T_{FCFS} = \frac{\lambda(\overline{s^2} + 2c_1\sqrt{A}\bar{s} + c_2A)}{2(1 - \lambda c_1\sqrt{A} - \rho)} + \bar{s} + c_1\sqrt{A}. \tag{26}$$

The stability condition for this policy is $\rho + \lambda c_1\sqrt{A} < 1$; therefore, this policy is unstable for values of ρ approaching 1. For $\lambda \to 0$, the first term in (26) approaches zero. Likewise, the second term in Theorem 1 also approaches zero as $\lambda \to 0$. So for the light traffic case we have

$$\frac{T_{FCFS}}{T^*} \leq \frac{\bar{s} + c_1\sqrt{A}}{\bar{s} + c_3\sqrt{A}}, \quad \text{as } \lambda \to 0.$$

Since $\bar{s}$ could be arbitrarily small, the worst case relative performance for this policy in light traffic is $\frac{T_{FCFS}}{T^*} \leq c_1/c_3 \approx 1.36$.

The FCFS policy can be modified to yield asymptotically optimal performance in light traffic as follows: Consider the policy of locating the server at the median of $\mathcal{A}$ and following a FCFS policy, where the server travels directly to the service site from the median, services the customer, and then returns to the median after service is completed. We call this policy the stochastic queue median policy (SQM). As before, the server waits at the median if no customers are present in the system. Again, since locations are independent of the order of arrival and the number in queue, the system behaves as a $M/G/1$ queue; however, we have to be somewhat careful about counting travel time in this case. From a system viewpoint, each "service time" now includes the on-site service plus the *round trip* travel between the median and the service location. The system time of an individual customer, however, includes the wait in queue plus a *one way* travel to the service location plus the on-site service. Therefore, the average system time under this policy is given by

$$T_{SQM} = \frac{\lambda(\overline{s^2} + 4c_3\sqrt{A}\bar{s} + 4c_4A)}{2(1 - 2\lambda c_3\sqrt{A} - \rho)} + \bar{s} + c_3\sqrt{A}, \tag{27}$$

where $c_3 \approx 0.383, c_4 = \frac{1}{6}$. The stability condition for this policy is $2\lambda c_3\sqrt{A} + \rho < 1$.

Letting λ approach zero, the first term above goes to zero and since c_3 is the constant

of the lower bound in Theorem 1 we get

$$\frac{T_{SQM}}{T^*} \to 1, \quad \text{as } \lambda \to 0. \tag{28}$$

This argument can be generalized to arbitrary regions $\mathcal{A}$ by substituting $E[\|X - x^*\|]$ for $c_3\sqrt{A}$ and $E[\|X - x^*\|^2]$ for $c_4 A$ in (27). Therefore, we have the following theorem.

Theorem 3 *The SQM policy of locating the server at the median of the region $\mathcal{A}$ and servicing customers in a FCFS order (returning to the median after each service is completed) is asymptotically optimal for the DTRP as λ approaches zero.*

This is an intuitively satisfying if not altogether surprising result. It is conjectured by Psaraftis in [25]. It is also analogous to the results achieved by Berman et. al. [4] and Batta et. al. [2] for the optimal location of a server on a network operated under a FCFS policy. Our result is somewhat stronger because our lower bound is on *all* policies, not just FCFS policies. Therefore it establishes not only the optimality of the median location for the SQM, but also the optimality of the SQM discipline itself.

The FCFS and SQM policies become unstable for $\rho \to 1$. The reason for this is that the average distance traveled per service, $\bar{d}$, remains fixed, yet the stability condition (22) implies $\bar{d} < \frac{1-\rho}{\lambda}$, so $\bar{d}$ must decrease as ρ (and λ) are increased. As shown below, a policy that is stable for all values of ρ must increasingly restrict the distance the server is willing to travel between services as the traffic intensity increases.

4.2 The Partitioning Policy

In this section we examine a policy that achieves the restriction on $\bar{d}$ mentioned above through a partition of the service region $\mathcal{A}$. The analysis relies on results for symmetric, cyclic queues, so readers unfamiliar with this area are encouraged to reexamine the definitions and results in Section 1.

Consider the following policy for the DTRP, which we call PART: The square region $\mathcal{A}$ is divided into m^2 subregions, where $m > 1$ is a given integer that parameterizes the policy. Within each subregion, customers are served using a FCFS discipline identical to the first FCFS policy of the previous section. The server services a subregion until their are no more customers left in that subregion. It then moves on to the next subregion and services it until no more customers are left, etc. The sequence of regions the server follows is shown in

17

<- Direction of Travel

Figure 1: Sequence for Serving Subregions PART Policy ($m = 4$)

Figure 2: PART Projection Policy for Moving to Adjacent Subregion

Figure 1 for the case $m = 4$. (Note that the server always moves to an *adjacent* subregion.) The pattern is continuously repeated.

To move from one subregion to the next, the server uses the projection rule shown in Figure 2. Its last location in a given subregion is simply "projected" onto the next subregion to determine the server's new starting location. The server then travels in a straight line between these two locations. As a result of this rule, note that the distance traveled between subregions is a constant $\frac{\sqrt{A}}{m}$, and that each starting location is uniformly distributed and independent of the locations of customers in the new subregion. These properties of the starting location simplify the analysis. In practice, one might use a more intelligent rule such as moving directly to the first customer in the next subregion. The total travel distance of this tour is $m^2(\sqrt{A}/m) = m\sqrt{A}$.

18

Notice that to construct the pattern shown in Figure 1, m must be even. If m is odd, the server ends up in the upper right subregion and must travel to the lower right subregion to restart the cycle. This adds an additional $\sqrt{A} - \sqrt{A}/m$ to the total travel distance. To simplify the analysis, we use only the expression for even m. As shown below, m must be large in heavy traffic, so for $\rho \to 1$ the error in total travel distance is negligible.

Each subregion behaves as an $M/G/1$ queue with an arrival rate of $\frac{\lambda}{m^2}$, and first and second moments of $\bar{s} + c_1 \frac{\sqrt{A}}{m}$ and $\overline{s^2} + 2c_1 \bar{s} \frac{\sqrt{A}}{m} + c_2 \frac{A}{m^2}$ respectively ($c_1 \approx 0.52, c_2 = \frac{1}{3}$). The policy as a whole behaves as a cyclic queue with $k = m^2$ queues and exhaustive service, where the total travel time around the cycle is $m\sqrt{A}$ and the queue parameters are those given above. Again, as with the FCFS policy, the travel times are not mutually independent. However, they are identically distributed and independent of the number in queue. Therefore, the analysis in [5] still holds. Recalling that the expression in (3) is for the waiting time in queue only, the average system time for this policy is given by

$$T_{PART} = \frac{\lambda(\overline{s^2} + 2c_1 \bar{s} \frac{\sqrt{A}}{m} + c_2 \frac{A}{m^2})}{2(1 - \lambda(\bar{s} + c_1 \frac{\sqrt{A}}{m}))} + \frac{1 - \frac{\lambda}{m^2}(\bar{s} + c_1 \frac{\sqrt{A}}{m})}{2(1 - \lambda(\bar{s} + c_1 \frac{\sqrt{A}}{m}))} m\sqrt{A} + c_1 \frac{\sqrt{A}}{m} + \bar{s}. \qquad (29)$$

The stability condition is

$$\lambda(\bar{s} + c_1 \frac{\sqrt{A}}{m}) < 1 \qquad \Leftrightarrow \qquad m > \frac{c_1 \lambda \sqrt{A}}{1 - \rho}.$$

Defining the critical value m_c by

$$m_c \equiv \frac{c_1 \lambda \sqrt{A}}{1 - \rho}, \qquad (30)$$

the stability conditions requires $m > m_c$. Note that for any $\rho < 1$ we can find an $m > m_c$ such that this policy is stable. Since the optimal policy has a waiting time no greater than the PART policy, we have the following theorem:

Theorem 4 *There exists a DTRP policy that has a finite waiting time for all $\rho < 1$ (the PART policy) and hence there exists an optimal policy for all $\rho < 1$.*

Note that this establishes $\rho < 1$ as the stability condition for the DTRP. It also shows that the size of the service region does not place a limit on the traffic intensity, which is rather surprising.

For given system parameters λ, $\bar{s}$, $\overline{s^2}$ and A, one could perform a one dimensional optimization over $m > 1$ using (29) to get the optimum number of partitions; however,

since equation (29) is quite complicated, we concentrate on finding the optimal value, m^*, for the heavy traffic case.

For $(\rho \to 1)$, (30) implies that any feasible m is large $(m > m_c)$. Therefore ignoring the $O(1/m)$ and smaller terms in the numerators of (29) we obtain

$$T_{PART} \approx \frac{\lambda \overline{s^2} + m\sqrt{A}}{2(1 - \rho - \lambda c_1 \sqrt{A}/m)} = \frac{m^2 \sqrt{A} + m\lambda \overline{s^2}}{2(m(1 - \rho) - \lambda c_1 \sqrt{A})}. \tag{31}$$

Differentiating the above with respect to m and setting the result equal to zero, we get the following critical points

$$\frac{\lambda c_1 \sqrt{A} \pm \sqrt{\lambda^2 c_1^2 A + (1 - \rho)\lambda^2 c_1 \overline{s^2}}}{1 - \rho}.$$

Only the positive root is feasible. For $\rho \to 1$ the second term under the radical approaches zero; therefore

$$m^* \approx \frac{2\lambda c_1 \sqrt{A}}{1 - \rho} = 2m_c.$$

If we substitute this value into (31), the optimal waiting time in heavy traffic is

$$T_{PART} \approx 2c_1 \frac{\lambda A}{(1 - \rho)^2} + \frac{\lambda \overline{s^2}}{1 - \rho}. \tag{32}$$

For $\rho \to 1$, the first term above dominates; therefore (recalling the bound in Theorem 2) we have

$$\frac{T_{PART1}}{T^*} \leq \frac{2c_1}{\gamma^2}, \quad \text{as } \rho \to 1. \tag{33}$$

This says that the PART system time is within a constant factor of the optimum in heavy traffic, though the provable factor is indeed quite large (about 44). If the conjectured value of $\gamma = 1/2$ is used, the factor is a more reasonable 4.2.

4.3 The Traveling Salesman Policy

The traveling salesman policy (TSP for short) is based on collecting customers into sets that can then be served using an optimal TSP tour. Let $\mathcal{N}_k$ denote the kth set of n customers to arrive, where n is a given constant that parameterizes the policy, e.g. $\mathcal{N}_1$ is the set of customers $1, \ldots, n$, $\mathcal{N}_2$ is the set of customers $n + 1, \ldots, 2n$, etc. Assume the server operates out of a depot at a random location in $\mathcal{A}$. When all customers in set $\mathcal{N}_1$ have arrived, we form a TSP tour of these customers starting and ending at the depot. Customers are then serviced by following the tour. If all $\mathcal{N}_2$ customers have arrived when the tour of

$\mathcal{N}_1$ is completed, they are serviced using a TSP tour; otherwise, the server waits until all $\mathcal{N}_2$ customers arrive before serving it. In this manner, sets are serviced in a FCFS order. Observe also that queueing of sets can occur.

Suppose one considers the set $\mathcal{N}_k$ to be the *kth* customer. Since the interarrival time (time for n new demands to arrive) and service time (n on-sites services plus the travel time around the tour) of sets are i.i.d., the service of sets forms a $GI/G/1$ queue, where the interarrival distribution is Erlang of order n. The mean and variance of the interarrival times for sets are n/λ and n/λ^2 respectively. The service time of sets is the sum of the travel time around the tour, which we denote L_n, and the n on-site service times. If we let $E[L_n]$ and $var[L_n]$ denote, respectively, the mean and variance of L_n, then the expected value of the service time of a set is $E[L_n] + n\bar{s}$ and the variance is $var(L_n) + n\sigma_s^2$, where $\sigma_s^2 = \overline{s^2} - \bar{s}^2$ is the variance of the on-site service time.

We are now in a position to apply the $GI/G/1$ upper bound (1) for the average waiting time of sets, W_{set}. This gives

$$W_{set} \leq \frac{\frac{\lambda}{n}(\frac{n}{\lambda^2} + var[L_n] + n\sigma_s^2)}{2(1 - \frac{\lambda}{n}(E[L_n] + n\bar{s}))} \tag{34}$$

$$= \frac{\lambda(1/\lambda^2 + \frac{var[L_n]}{n} + \sigma_s^2)}{2(1 - \rho - \lambda\frac{E[L_n]}{n})}. \tag{35}$$

As we show below, in order for the policy to be stable in heavy traffic n has to be large. Thus, because the locations of points are uniform and i.i.d. in the region, we have from the asymptotic results for the TSP (10) and (11) that

$$\frac{E[L_n]}{n} \approx \beta_{TSP}\frac{\sqrt{A}}{\sqrt{n}}, \tag{36}$$

and

$$\frac{var[L_n]}{n} \approx 0, \tag{37}$$

where the approximations become exact for $n \to \infty$. In order to simplify the final expressions, we have neglected the difference between $n + 1$ and n in the above expressions. (The tour includes n points plus the depot.) Since n is large, the difference is negligible. Therefore, for large n

$$W_{set} \leq \frac{\lambda(1/\lambda^2 + \sigma_s^2)}{2(1 - \rho - \lambda\beta_{TSP}\frac{\sqrt{A}}{\sqrt{n}})}. \tag{38}$$

For stability, we require $\rho + \lambda \beta_{TSP} \frac{\sqrt{A}}{\sqrt{n}} < 1$, which implies

$$n > \frac{\lambda^2 \beta_{TSP}^2 A}{(1-\rho)^2}. \tag{39}$$

For $\rho \to 1$ n *must* be large, and thus using asymptotic TSP results is indeed justified.

The waiting time given in (38) is not itself an upper bound on the wait for service of an individual customer; 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 customer in that set. Therefore, we must add to (38) the time a customer waits for its set to form and also the time it takes to complete service of the customer once the customer's set enters service. By conditioning on the position that a given customer takes within its set, it is easy to show that the average wait for a customer's set to form is $\frac{n-1}{2\lambda} \le \frac{n}{2\lambda}$. By doing the same conditioning and noting that the travel time around the tour is no more that the length of the tour itself, the expected wait for service once a customer's set enters service is no more than $\beta_{TSP}\sqrt{nA} + \frac{1}{n}\sum_{k=1}^{N} k\bar{s} \le \beta_{TSP}\sqrt{nA} + \frac{n}{2}\bar{s}$. Therefore, if the total system time is denoted T_{TSP},

$$T_{TSP} \le \frac{\lambda(1/\lambda^2 + \sigma_s^2)}{2(1-\rho-\lambda\beta_{TSP}\frac{\sqrt{A}}{\sqrt{n}})} + \frac{n(1+\rho)}{2\lambda} + \beta_{TSP}\sqrt{nA}. \tag{40}$$

We would like to minimize (40) with respect to n to get the least upper bound. (One can verify that (40) is convex, so there is indeed a minimum.) First, however, consider a change of variable

$$y = \frac{\lambda\beta\sqrt{A}}{(1-\rho)\sqrt{n}}.$$

Physically, y represents is ratio of the average distance, $\bar{d} = \frac{\beta_{TSP}\sqrt{A}}{\sqrt{n}}$ to its critical value $\frac{1-\rho}{\lambda}$. With this change,

$$T_{TSP} \le \frac{\lambda(1/\lambda^2 + \sigma_s^2)}{2(1-\rho)(1-y)} + \frac{\lambda\beta_{TSP}^2 A(1+\rho)}{2(1-\rho)^2 y^2} + \frac{\lambda\beta_{TSP}^2 A}{(1-\rho)y}. \tag{41}$$

For $\rho \to 1$, one can verify that the optimum y approaches 1. Therefore, by linearizing the last two terms above about $y = 1$, an approximate optimum value, y^*, is

$$y^* \approx 1 - \frac{\sqrt{(1/\lambda^2 + \sigma_s^2)(1-\rho)}}{2\beta_{TSP}\sqrt{A}}.$$

Substituting this approximation into (41) and noting that for $\rho \to 1$ the approximate y^* approaches 1 we have

$$T_{TSP} \le \beta_{TSP}^2 \frac{\lambda A}{(1-\rho)^2} + \frac{\beta_{TSP}\lambda\sqrt{A(1/\lambda^2 + \sigma_s^2)}}{(1-\rho)^{3/2}} + \frac{\beta_{TSP}^2\lambda A}{(1-\rho)} \qquad \rho \to 1.$$

Again, the leading term is proportional to $\frac{\lambda A}{(1-\rho)^3}$. Therefore, using Theorem 2

$$\frac{T_{TSP}}{T^*} \le \frac{\beta^2_{TSP}}{\gamma^2} \quad \rho \to 1.$$

The best estimate to date of β_{TSP} is approximately 0.72 [14], so the TSP policy has a system time in heavy traffic about half that of the partitioning policy. (In practice, heuristic rather than optimal tours would be used to reduce the computational burden, which would produce slightly higher system times.) These results suggest that the policy of forming successive TSP tours, which is a reasonable policy in practice, is quite good theoretically. In addition to providing a theoretical guarantee, the results give a practical means of optimally sizing routes for such policies by either minimizing the right hand side of (40) or using the approximate y^*.

4.4 The Space Filling Curve Policy

We next analyze a policy based on space filling curves which we call the SFC policy. It was first proposed by Bartholdi and Platzman [1]. The reader is encouraged to reexamine Section 1 for notation and basic results related to space filling curves. Let C and ψ be defined as in Section 1, and let the DTRP service region, $\mathcal{A}$, be a square of area A. Suppose we maintain the preimages of all customer in the system (i.e. their positions in C). Then the SFC policy is to service customers as they are encountered in repeated clockwise sweeps of the circle C. (Note that one could treat a depot as a permanent "customer" and visit it once per sweep.)

We now analyze this policy. Consider a randomly tagged arrival and let W_0 denote the waiting time of the tagged arrival, $\mathcal{N}_0$ denote the set of locations of the $N_0 \equiv |\mathcal{N}_0|$ customers served prior to the tagged customer, and L denote the length of the path from the server's location through the points in $\mathcal{N}_0$ to the tagged customer's location which is induced by the SPC rule. Finally, let s_i be the on-site service time of customer $i \in \mathcal{N}_0$, and R be the residual service time of the customer under service. Then

$$W_0 = \sum_{i=1}^{N_0} s_i + L + R.$$

Taking expectation on both sides gives

$$W = E[N_0]\bar{s} + E[L] + \frac{\lambda \overline{s^2}}{2}. \tag{42}$$

23

Since in steady state the expected number of customers served during a wait equals the expected number who arrive, $E[N_0] = N = \lambda W$. Also, since L is the length of a path through $N_0 + 2$ points in $\mathcal{A}$, $L \leq 2\sqrt{(N_0 + 2)A}$. Therefore,

$$
\begin{aligned}
E[L] &\leq 2E[\sqrt{(N_0 + 2)A}] && (43) \\
&\leq 2\sqrt{(N + 2)A} && (Jensen's Ineq.) \\
&\leq 2\sqrt{\lambda W A} + 2\sqrt{2A}.
\end{aligned}
$$

Substituting these results into (42) we obtain the following quadratic inequality:

$$
W - \frac{2\sqrt{\lambda A}}{1 - \rho}\sqrt{W} - \frac{\lambda \overline{s^2} + 4\sqrt{2A}}{2(1 - \rho)} \leq 0.
$$

Solving the above for W and recalling that $T = W + \overline{s}$ it is straightforward to show that

$$
T_{SFC} \leq \gamma_{SFC}^2 \frac{\lambda A}{(1 - \rho)^2} + o((1 - \rho)^{-2}), \tag{44}
$$

where $\gamma_{SFC} \leq 2$ and $o((1 - \rho)^{-2})$ denotes terms that increase more slowly than $(1 - \rho)^{-2}$ as $\rho \to 1$. This shows the SPC policy is within a constant factor of optimal. The constant γ_{SFC} obtained by the above argument, however, is based on worst-case tours and is probably too large. If one assumes that the clockwise interval between the preimages of the server and the customer is a uniform $[0, 1]$ random variable and that the $\mathcal{N}_0$ points are approximately uniformly distributed on this interval, then a constant of $\gamma_{SFC} \approx \frac{2}{3}\beta_{SPC} \approx 0.64$ is obtained.

To estimate γ_{SFC} more precisely, we performed simulation experiments. The method of *batch means* (see [20]) was used to estimate the steady state value of T_{SFC}. In this method, customers are grouped into batches of a fixed size. If the batch size is large enough, the sample means from each batch are approximately uncorrelated and normally distributed [21]. (We used 200 times the minimum average number in the system given by Theorem 2 as our batch size.) The sample mean and variance of the individual batch means were then used in a t-test to estimate T_{SFC}. The simulation was terminated when the 99% confidence interval about the estimate reached a width less than 10% of the value of the estimate. The batch means method was selected because the busy periods of the SFC policy were quite long (indeed, almost nonterminating) at high utilization values, which precluded the use of techniques based on regeneration points.

The simulation was run for $A = 1$ and a range of parameter values ρ, $\overline{s}$ and $\overline{s^2}$. Figure 3 shows one example of the simulation estimate of T_{SFC} plotted against $\lambda A/(1 - \rho)^2$ for the

```
                                    SFC
                               Slope = 0.440
    100

                                           NN
     80                                 Slope = 0.406


  E  60
  Φ
  E
  Φ
  n
  >  40


     20


      0
         0      100      200      300
                  λA/(1-ρ)2
```
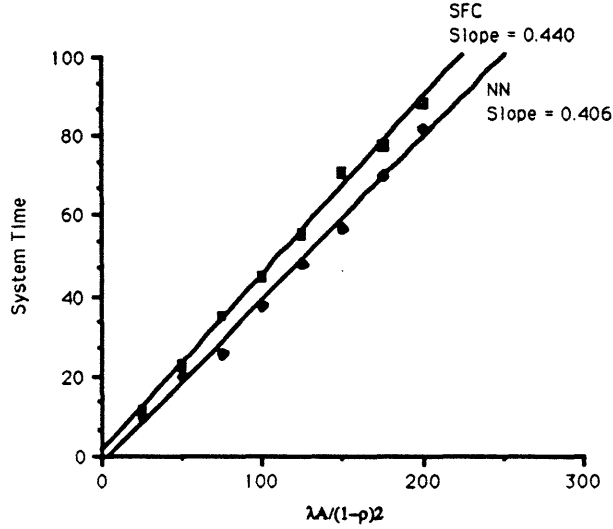
Figure 3: Simulation Results: T_{SFC} and T_{NN} vs. $\lambda A/(1 - \rho)^2$

case $A = 1$, $\bar{s} = 0.1$ and $\overline{s^2} = 0.01$ (zero variance). Each point is a different value of ρ in the range $0.5 - 0.8$ The results showed that γ_{SFC} is approximately 0.66, which is very close to the approximate value of $\frac{2}{3}\beta_{SPC}$. The system time for this policy is therefore about 15% lower than that of the TSP policy. It is also much more computationally efficient.

4.5 The Nearest Neighbor Policy

The last policy we consider is to serve the closest available customer after every service completion (nearest neighbor (NN) policy). The motivations for considering such a policy are: 1) the nearest neighbor was used in the heavy traffic lower bound of Theorem 2, and 2) the shortest processing time (SPT) rule is known to be optimal for the classical $M/G/1$ queue [10]. As mentioned before, however, the travel component of service times in the DTRP depends on the service sequence, so the classical $M/G/1$ results are not directly applicable; they are only suggestive.

Because of the dependencies among the travel distances d_i, we were unable to obtain rigorous analytical results for the NN policy. However, if one assumes there exists a constant γ_{NN} such that

$$E[d_i | N_T] \le \gamma_{NN} \frac{\sqrt{A}}{\sqrt{N_T}}, \qquad (45)$$

where N_T is the number of customers in the system at a completion epoch, then by using

a modification of the argument in [18] Section 5.5, it is possible to show that

$$T_{NN} \leq \gamma_{NN}^2 \frac{\lambda A}{(1-\rho)^2} \qquad \rho \to 1,$$

where T_{NN} denotes the system time of the NN policy. The assumption (45) is analogous to (21) but unlike (21) has not been established formally.

We therefore performed simulation experiments identical to those for the SFC policy to verify the asymptotic behavior of T_{NN} and estimate γ_{NN}. The results showed that γ_{NN} is approximately 0.64. (See Figure 3.) This means that T_{NN} is about 10% lower than T_{SFC} and about 20% lower than T_{TSP}.

The results again confirmed that the system time T_{NN} follows the $\frac{\lambda A}{(1-\rho)^2}$ growth predicted by the lower bound in Theorem 2. Figure 3 clearly shows this highly linear relationship.

4.6 A Numerical Example

To illustrate the relative performance of the various DTRP policies, the system time of each policy was calculated (simulated in the case of SFC and NN policies) for the case $A = 1$, $\bar{s} = 0.1$ and $\overline{s^2} = 0.01$ (zero variance) for a range of values of ρ. For the parameterized policies (PART and TSP), numerical optimization was performed to find the best parameter for each value of ρ. The results showed that the FCFS, SQM, SFC and NN policies performed well in light traffic but the FCFS and SQM policies were unstable for $\rho > 0.2$. The PART, TSP, SFC and NN policies performed best in heavy traffic. Results for each group are graphed separately.

Figure 4 shows system times as a function of ρ for the light traffic case. The lower bound is also included. Note that although the SQM policy is asymptotically optimal as $\rho \to 0$, it is quickly surpassed by the FCFS policy as ρ increases. This is due to the extra travel distance of the SQM policy, which hinders the policy as queueing sets in. Also note that both policies reach their saturation points for relatively low values of ρ. The SFC and NN policies were comparable to the FCFS policy in very light traffic, which is to be expected since they essentially behave like the FCFS policy in this case. For $\rho > 0.05$ the SFC and NN policies quickly surpass the FCFC and SQM policies. Notice that the NN policy consistently performed better than the SFC policy even in the light traffic cases.

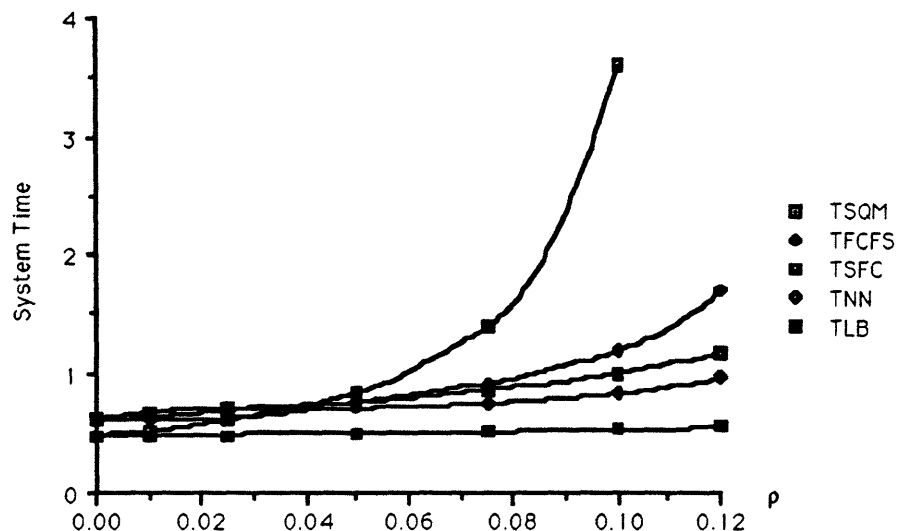The heavy traffic results are shown in Figure 5. Note that the curves have nearly identical

Figure 4: System Times for Light Traffic Case: Numerical Example

shapes as one would expect from the $\frac{\lambda A}{(1-\rho)^2}$ asymptotic behavior of each policy. (Only the constant of proportionality differs.) The graphs show the sharp increase in system time as the traffic intensity increases. The NN policy is the best in this case with the SFC a close second best. The TSP and especially the PART policy are less effective.

This example suggests that both the SFC and NN policies are effective over a wide range of traffic intensities. Indeed, if one locates a depot at the median of the region $\mathcal{A}$ and treats it as a permanent "customer", then both these policies can be made to behave like the SQM policy as $\rho \rightarrow 1$. These policies also have the advantage of being nonparametric (i.e. the system parameters are not needed to implement them as is the case for the TSP and PART policies) and are therefore self regulating. This feature is especially desirable for system that operate under highly variable and unpredictable traffic conditions.
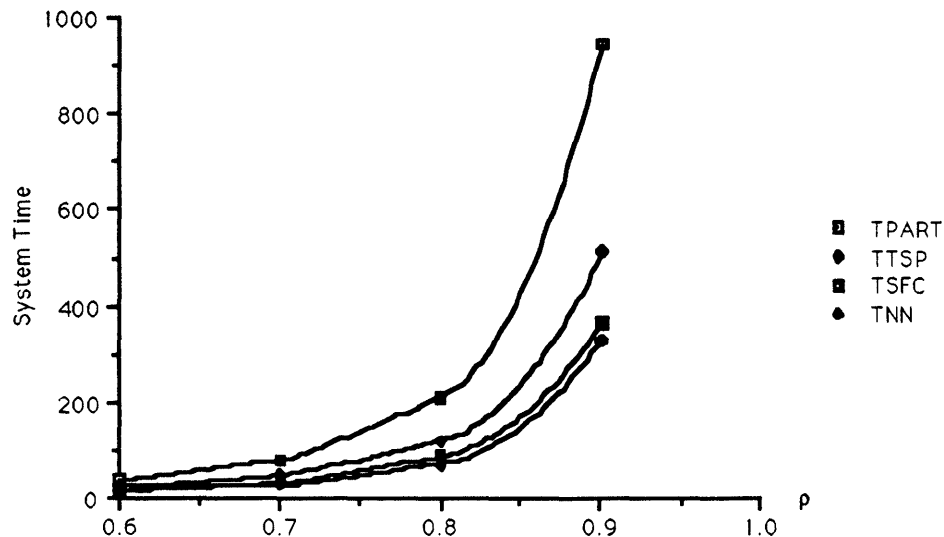
Figure 5: System Time for Heavy Traffic Policies: Numerical Example

5   Concluding Remarks

We presented a new model for dynamic vehicle routing problems that attempts to capture the dynamic and stochastic environment in which real-world systems operate. It constitutes a major departure from traditional static and deterministic models. Several application areas were suggested for which this model is appropriate. We derived lower bounds on the optimal system time and characterized the performance of several diverse policies.

The stochastic queue median policy, in which we strategically locate a depot and then follow a FCFS service order, was shown to be optimal in light traffic. As the traffic intensity increases, however, FCFS policies become unstable. We then showed that the partitioning policy behaved reasonably well in heavy traffic since it has a constant factor performance guarantee and finite system time for all values of $\rho < 1$.

In heavy traffic, the best policies were the TSP, SFC and NN. The SFC and NN policies have a desirable self regulating behavior, while the TSP policy has the advantage of returning regularly to the depot. The TSP and SFC would appear to be more "fair" than the NN policy since they partially obey a FCFS discipline (i.e. sets are served in FCFS order in the

28

case of the TSP policy and for the SFC policy, the entire region is periodically "swept" by the server). In addition, they have provable performance guarantees. The NN policy, on the other hand, has system times about 10% lower than the SFC policy and 20% lower than the TSP strategy according to our simulation study. It does not, however, have a provable performance guarantee.

These policies, though quite diverse, have identical asymptotic behavior in heavy traffic. Their asymptotic system time is proportional to $(1 - \rho)^{-2}$ and does not depend on the service time variation $(\overline{s^2})$. This is in stark contrast to the behavior of traditional queues, and it illustrates the unique insights that can be obtained by considering combined queueing/routing models.

We believe that this class of dynamic vehicle routing problems constitutes a very interesting and realistic class of models, and as such deserves additional attention. An obvious extension is to multiple server (m-vehicle) models. This is a topic we have recently investigated in [9] where similar bounds and policies are established. In particular, the system time is shown to have a $\frac{\lambda A}{m^2(1-\rho)^2}$ behavior in heavy traffic. One might also investigate the effect of vehicle capacity. Our preliminary results here suggest that the stability condition is no longer independent of the service region size in the capacitated case. Finally, one could certainly construct other DTRP policies and analyze them using the techniques of Section 1.

Acknowledgements

References

[1] Bartholdi, J.J. and Platzman, L.K. (1988), "Heuristics Based on Spacefilling Curves for Combinatorial Problems in Euclidean Space", *Mgmt. Sci.*, 34, 291-305.

[2] Batta, R., R.C. Larson and A.R. Odoni (1988), "A Single Server Priority Queueing Location Model", *Networks* 18, 87-103.

[3] Beardwood J., Halton J. and Hammersley J. (1959), "The Shortest Path Through Many Points", *Proc. Camb. Phil. Soc.*, 55, 299-327.

[4] Berman,O., S.S. Chiu, R.C. Larson, A.R. Odoni and R. Batta (1989), "Location of Mobile Units in a Stochastic Environment", in *Discrete Location Theory* (P.B. Mirchandani and R.L. Francis, eds.) Wiley, New York, in press.

[5] Bertsekas, D. and R. Gallager (1987), *Data Networks*, Prentice Hall, Englewood Cliffs.

[6] Bertsimas, D. (1988), "The Probabilistic Vehicle Routing Problem", MIT Sloan School of Management Working Paper No. 2067-88.

[7] Bertsimas, D. (1988), "Probabilistic Combinatorial Optimization Problems", Operations Research Center, MIT, Cambridge, MA, Technical Report No. 194.

[8] Bertsimas, D., Jaillet, P. and Odoni, A. (1988) "A Priori Optimization", December 1988, submitted for publication to *Operations Research*.

[9] Bertsimas, D. and van Ryzin, G. (1990), "The m-Server Dynamic Traveling Repairman Problem: Bounds, Policies and Optimization Problems", in preparation.

[10] Conway, R. W., Maxwell, W. L., and Miller, L. W., *Theory of Scheduling*, Addison-Wesley, Reading, Mass., 1967.

[11] Ferguson M. and Aminetzah Y. (1985), "Exact Results for Non-Symmetric Token Ring Systems", *IEEE Transactions on Communication*, COM-33(3), 223-231.

[12] Haimovich, M. and Magnanti, T.L. (1988), "Extremum Properties of Hexagonal Partitioning and the Uniform Distribution in Euclidean Location", *SIAM J. Disc. Math.*, 1, 50-64.

[13] Jaillet, P. (1988) "A Priori Solution of a Traveling Salesman Problem in Which a Random Subset of the Customers Are Visited", *Operations Research* 36:929-936.

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

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

[16] Kingman, J.F.C. (1962), "Some Inequalities for the Queue $GI/G/1$", *Biometrika*, 383-392.

[17] Kleinrock, L. (1976), *Queueing Systems, Vol 1: Theory*, Wiley, New York.

[18] Kleinrock, L. (1976), *Queueing Systems, Vol 2: Computer Applications*, Wiley, New York.

[19] Larson, R. and Odoni, A. (1981), *Urban Operations Research*, Prentice Hall, Englewood Cliffs.

[20] Law, A.M. and Kelton, W.D. (1982), *Simulation Modeling and Analysis*, McGraw-Hill, New York.

[21] Law, A.M. and Carson, J.S. (1979), "A Sequential Procedure for Determining the Length of a Steady-State Simulation", *Oper. Res.*, 27, 1011-1025.

[22] 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.

[23] Minieka E.,"The Delivery Man Problem on a Tree Network", unpublished manuscript.

[24] L.K. Platzman and J.J. Bartholdi (1983), "Spacefilling Curves and the Planar Travelling Salesman Problem", PDRC Technical Report 83-02, Georgia Institute of Technology, 1983. (To appear in *JACM*)

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

[26] Ross, S.M., (1983) *Stochastic Processes*, Wiley, New York.

[27] Sahni, S. and Gonzalez, T. (1976), "P-complete Approximation Problems", *Journal of ACM*, 23, 555-565.

[28] Steele, J.M. (1981), "Subadditive Euclidean Functionals and Nonlinear Growth in Geometric Probability", *Ann. Prob.* 9, 365-376.

[29] Steele, J.M. (1988), "Seedlings in the Theory of Shortest Paths", Princeton University Technical Report.