The Dynamic Traveling Repairman Problem

Dimitris Bertsimas
Sloan School of Management
Massachusetts Institute of Technology

Garrett van Ryzin
Operations Research Center
Massachusetts Institute of Technology

MIT Sloan School Working Paper No. 3036-89-MS

The Dynamic Traveling Repairman Problem

Dimitris J. Bertsimas * Garrett van Ryzin †

July 10, 1989

Abstract

We propose and analyze a generic mathematical model for dynamic 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 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.

Key words. Dynamic routing, repairman and traveling salesman problem, queueing, 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 ECS-8717970.

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 in large part to the problem's richness and inherent elegance; however, its popularity is also due to the fact that the TSP is encountered frequently in practical problems, both directly and as a subproblem. Yet, in many of these practical applications, the TSP is a deterministic, static approximation to a problem which is, in reality, both probabilistic and time varying (dynamic). Also, there are often costs associated with intervisit delays that are not captured in the objective of minimizing travel distance.

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 (delivery) costs. In a real system, however, the demands may arrive randomly in time and the dispatching of vehicles may be 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. The primary applications that motivate our investigation 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:

1. The demands from the remote locations are requests for replenishment of stock that occur randomly in time. In this case, large waiting times mean that large inventories are needed at the remote sites. Thus, decreasing the wait for delivery is important for reducing inventories.

2. 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.

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

4. In sequencing airplanes for landing in major airports it might be more appropriate to

minimize the total average waiting time in the air rather than minimizing the time the last airplane lands, which is exactly the TSP objective. In this example demand for service is dynamically changing. A similar application is the problem of managing a fleet of taxis to minimize the average waiting time of customers.

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 average amount of time each prospective customer spends contemplating his purchase!

Motivated by these application areas, we propose and analyze a generic mathematical model which we call the dynamic traveling repairman problem (DTRP). The DTRP is defined as follows: a region $\mathcal{A}$ contains a vehicle (server) that travels at unit velocity. Demands for service arrive according to a Poisson process and, upon arrival, are independently and uniformly assigned a location within $\mathcal{A}$. The problem is to find a policy for servicing the demands that minimizes the average system time, which includes both the waiting time in queue and the on-site service time.

The DTRP 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.

Indeed, Psaraftis [21] provides an extensive list of other characteristics that distinguish dynamic versions of the TSP from their static counterparts and points out that, in general, little is known about dynamic vehicle routing problems. More importantly, Psaraftis [21] defines the dynamic traveling salesman problem (DTSP), which motivated our investigation on dynamic vehicle routing problems, as follows. 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 time t_{ij} (which can be stochastic) 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). In comparison with the DTSP, the DTRP is defined in the Euclidean plane and specifically optimizes over the total system time.

Some of the above characteristics have been considered before in isolation in the literature. The first important characteristic of the DTRP is that the objective is to minimize waiting 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 [22]). In contrast with the TSP, which is trivial on trees, the TRP seems difficult on trees. Minieka [20] 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 problem.

Jaillet [10], Bertsimas [5] and Bertsimas, Jaillet and Odoni [6] address the first and third characteristic 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) which are defined 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 case the vehicle has capacity Q then 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, i.e. it is solved a priori using only the probabilistic information.

A distinct characteristic of the DTRP is that it incorporates queueing phenomena into the routing problem. Queueing considerations in the context of location problems have been considered in Berman et. al. [3] and Batta et. al. [1]. 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. In particular, this model, which is very appropriate for the location of emergency servers, assumes that arrivals occur in a dynamic manner according to a Poisson process. A vehicle (server) is dispatched from a central depot and then returns to the depot again. The problem is to locate the depot on a network so that the mean queueing delay and mean travel time is minimized. In our setting, the SQMP can be seen as a particular case of the DTRP in which the policy followed is to strategically locate the server and then follow a FCFS dispatching rule. We compare the performance ot this policy with other policies in Section 4.1.

The DTRP in the Euclidean plane is analyzed using a variety of techniques from combinatorial optimization, queueing theory, geometrical probability and simulation. Our strategy is the following: first, we establish some lower bounds on the average system time for all policies. Then, we propose several policies and find exact analytic expressions for their performance. 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 paper is organized as follows. Since we use a variety of results from several areas, we briefly describe these results 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 (FCFS, stochastic queue median, partitioning, TSP and nearest neighbor policies). In this section, we prove that the stochastic queue median policy is optimal in the light traffic limit and that several policies are within a constant factor of optimality in the heavy traffic limit. The policy with the best provable performance guarantee in heavy traffic is one based on forming TSP tours. It has an average system time of no more than 2.1 times the optimal system time. The nearest neighbor policy is shown via simulation to be within approximately 1.6 of the optimal system time, but this bound is not proven analytically. In Section 4.5, 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. Note that the average system time T is simply $W + \bar{s}$. Kingman [13] (see also Kleinrock [15]) 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 [15]) 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 to 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 either the exhaustive service policy, i.e. servicing each queue i until the queue is empty before proceeding, or the gated policy, where only customers in queue at the time of the servers arrival are served. The expected waiting time for this symmetric cyclic queue using an exhaustive service policy is given by (see Bertsekas and Gallager [4], p.156)

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

For the gated service policy, the waiting time is

$$W = \frac{\lambda \, \overline{s^2}}{2(1-\rho)} + \frac{(1+\frac{\rho}{k})}{2(1-\rho)} d. \tag{4}$$

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 [9]).

Jensen's Inequality

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

$$E[f(X)] \geq f(E[X]) \tag{5}$$

provided the expectations exist.

Geometrical Probability

Given n uniformly and independently distributed points $x_1, \ldots, x_n$ in a square of area A, then the following lower bound for the distance to the nearest neighbor is known (see for example Beardwood et. al. [2]):

$$E[\min_{j \neq i} |x_i - x_j|] \geq \frac{1}{2\sqrt{n}}\sqrt{A}. \tag{6}$$

Also,

$$E[|x_1 - x_2|] = c_1\sqrt{A}, \quad E[|x_1 - x_2|^2] = c_2 A, \tag{7}$$

where $c_1 \approx 0.52, c_2 = \frac{1}{3}$ (see Larson and Odoni [16], p.135).

If we let x^* denote the center of a square of area A, then it is known [16] that the first and second moment of the distance to a uniformly chosen point x are given by

$$E[|x^* - x|] = c_3\sqrt{A}, \quad E[|x^* - x|^2] = c_4 A, \tag{8}$$

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

$$\lim_{n\to\infty} \frac{L_n}{\sqrt{n}} = \beta_{TSP}\sqrt{A}. \tag{9}$$

with probability one (see [23], [19]). In his recent experimental work with very large scale TSP's, Johnson [11] estimated $\beta_{TSP} \approx 0.72$. In addition, it is also well known (see [19], p. 189) that $\lim_{n\to\infty} var(L_n) = O(1)$, and therefore

$$\lim_{n\to\infty} \frac{var(L_n)}{n} = 0. \tag{10}$$

2 Problem Definition and Notation

The DTRP was defined broadly in the introduction. We repeat the definition here for convenience and to establish 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. 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 $\overline{s^2}$. The fraction of time the server spends in on-site service is denoted ρ, and for stable systems $\rho = \lambda\bar{s}$. The system time of demand i, denoted T_i, is defined as the elapsed time between the arrival of the demand and the time the server completes the service of i. The problem is to find a policy for servicing demands that minimizes the steady-state, average system time, $T \equiv \lim_{i\to\infty} E[T_i]$.

We restrict the class of policies slightly by requiring that routing decisions be made only at service completion epochs. This means the server travels in straight lines between demand locations and is prohibited from making changes in "mid-course". We let d_i denote the straight-line travel time *to* the *ith* demand location from the location at which the server made the decision to service i. The term d_i is therefore 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 $\bar{d} \equiv \lim_{i\to\infty} E[d_i]$.

A final remark concerning the difference between the DTRP and the $M/G/1$ queue: in the DTRP, the total service time 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 times 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

To simplify the calculations and the presentation, it is assumed throughout the next two sections that the region $\mathcal{A}$ is a square of area A. This restriction can be relaxed without affecting the results, though the actual calculations may be more difficult. The strategy we follow is to first establish two simple but powerful lower bounds on the optimal expected system time, T^*. In Section 4, we 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

$$T = W^d + W^s + \bar{s}, \tag{11}$$

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

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 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 uniform location to a server location selected from this distribution. Now, suppose we had the option of locating the server in the best *a priori* location, x^*, prior to the arrival. 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

$$W^d \geq \min_{x_0 \in \mathcal{A}} E_x[||x - x_0||] \tag{12}$$

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 (8),

$$W^d \geq c_3\sqrt{A} \approx 0.383\sqrt{A} \tag{13}$$

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

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

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

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

By using (11) we obtain

$$W^s = \frac{\rho}{1-\rho}(W^d + \bar{s}) + \frac{\lambda \overline{s^2}}{2(1-\rho)}. \tag{14}$$

Combining (13) and (14) and noting that these bounds are true for *all* policies we get the first bound on T^*,

$$T^* \geq \frac{c_3\sqrt{A} + \bar{s}}{1-\rho} + \frac{\lambda \overline{s^2}}{2(1-\rho)}. \tag{15}$$

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

3.2 A Heavy Traffic Lower Bound

A second lower bound on T^* is obtained by examining the stability conditions for the DTRP. For the system to remain stable, the average amount of work (time) each customer requires from the server must not exceed the average time between arrivals,

$$\frac{1}{\lambda} > \bar{s} + \bar{d}. \tag{16}$$

Recall that $\bar{d}$ is the average travel time between customers and that the server makes service decisions at customer completion epochs.

Let d_i^* denote the distance between the server and the closest customer location at the epoch of the decision to service i; then by conditioning on the number of customers N in the system at the decision epoch we find

$$\bar{d} = E_N\left[E[d_i|N]\right] \geq E_N\left[E[d_i^*|N]\right] \tag{17}$$

Now, suppose the policy the server follows is such that the locations of the N customers in the system are independent and uniformly distributed. For example, this is the case for any policy in which service order is independent of customer locations, such as a FCFS, LCFS, or SIRO policy. In this case, $E[d_i^*|N]$ is bounded by (6),

$$E[d_i^*|N] \geq \frac{\sqrt{A}}{2\sqrt{N}} \tag{18}$$

For an arbitrary policy, the location of the N customers in the system at a completion epoch may not be uniformly distributed and independent; however, we make the assumption, that *for all stable policies* (18) is a valid lower bound. Indeed, we conjecture that this assumption is valid, though we have not been able to prove it. Arguing intuitively, the density of customers in the immediate vicinity of the server should be *lower* than the density of customers spread uniformly throughout the region due to the fact that recent service completions are, on average, close to the server's current location ($\overline{d}$ sufficiently small to satisfy (16)); thus, the customers are likely to be "thinned out" near the server. It is therefore reasonable that the expected value of d_i^* is as least as large as would be obtained if the N locations were uniform.

Assuming (18) is true generally, then by an identical argument to that of the $M/G/1$ queue, the number of customers in the system at service completion epochs has the same distribution as the number in the system at a random point in time; therefore, we can consider the random variable N above to be the number of customers in the system at a random instant. Applying Jensen's inequality to (18) and noting that $E[N] = \lambda T^*$ for the optimal policy, we have from (16) that

$$\frac{1}{\lambda} > \overline{s} + \frac{\sqrt{A}}{2\sqrt{\lambda T^*}},$$

which after rearrangement yields

$$T^* \geq \frac{\lambda A}{4(1-\rho)^2}. \tag{19}$$

This establishes the second lower bound. For $\lambda \to 0$ this bound approaches zero and is therefore not useful in the light traffic case. For $\rho \to 1$, on the other hand, this bound dominates (15).

The lower bound in (19) says that the waiting time for any policy must grow at least as fast as $(1-\rho)^{-2}$ rather than $(1-\rho)^{-1}$ as is the case for a classical queueing system. This is significantly different behavior from that of a traditional queueing system (e.g. the $M/G/1$ system). The reason for this difference lies in the geometry of the system; the bound (19) gives (via Little's Theorem) the minimum average number of customers that must be maintained in the system to ensure that the average travel distances, $\overline{d}$, satisfy the stability condition (16). This minimum number, however, grows much more rapidly than the average number in the system due simply to queueing delays. Indeed, as shown in the next section, *every* policy we analyzed that is stable for $\rho \to 1$ has a dominant term proportional to

$\lambda A/(1-\rho)^2$. This is further evidence that the lower bound (19) correctly describes the asymptotic behavior of the DTRP.

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 $\lambda \to 0$. These policies, however, are unstable for high utilizations; therefore, we turn next to a class of partitioning policies 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 are able to find exact expressions for the waiting time. Based on these explicit expressions, we determine the optimal number of partitions. These results show that the partitioning policies are within a constant factor of the lower bounds for all values of $\rho < 1$. They also establish $\rho < 1$ as the stability condition; that is, there exists stable policies for every $\rho < 1$. We then introduce a more sophisticated policy based on forming successive TSP tours. Its average system time is about 2.1 times the lower bound, which is nearly twice as good as the best partitioning policy. Finally, we examine the policy of serving the nearest neighbor. Because of analytical difficulties, we use simulation to analyze its performance. The simulation results show that this policy has an average system time within a factor of approximately 1.6 of the lower bound.

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.

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 they are identically distributed and also independent of the number in queue. Therefore, the Pollaczek-Khinchin (P-K) formula (2) still holds. (See [4] 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 service time are, by (7), $\bar{s} + c_1\sqrt{A}$ and $\overline{s^2} + 2c_1\sqrt{A}\bar{s} + c_2 A$ respectively, where $c_1 \approx 0.52, c_2 = \frac{1}{3}$. The average system time (waiting time plus service time) is therefore

$$T_{FCFS} = \frac{\lambda(\overline{s^2} + 2c_1\sqrt{A}\bar{s} + c_2 A)}{2(1 - \lambda c_1\sqrt{A} - \rho)} + \bar{s} + c_1\sqrt{A}. \tag{20}$$

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 (20) approaches zero. Likewise, the second term of (15) also approaches zero as $\lambda \to 0$; therefore, 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

$$T_{SQM} = \frac{\lambda(\overline{s^2} + 4c_3\sqrt{A}\bar{s} + 4c_4 A)}{2(1 - 2\lambda c_3\sqrt{A} - \rho)} + \bar{s} + c_3\sqrt{A}, \tag{21}$$

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 (15) we get

$$\frac{T_S}{T^*} \rightarrow 1, \qquad as\ \lambda \rightarrow 0. \tag{22}$$

This argument can be generalized to arbitrary regions $\mathcal{A}$; therefore, we have the following theorem.

Theorem 1 *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 analogous to the results achieved by Berman et. al. [3] and Batta et. al. [1] for the optimal location of a server on a network operated under a FCFS policy. Our result is somewhat stronger in that (15) is a lower bound on *all* policies, not just FCFS policies. It therefore establishes not only the optimality of the median location in light traffic, but also the optimality of the FCFS discipline itself; however, this is mitigated by the fact that very little queueing occurs in light traffic so the service discipline is insignificant compared to the server's location. As a practical matter, (15) can be used to bound the deviation from optimality of FCFS policies as λ is increased from zero.

Because of the stability conditions of the FCFS policies, T_{FCFS} becomes unbounded as $\rho + c_1\lambda\sqrt{A} \rightarrow 1$ and T_{SQM} becomes unbounded as $\rho + 2\lambda c_3\sqrt{A} \rightarrow 1$. As we show in the next subsection, there are stable policies for $\rho < 1$; therefore, the performance guarantee on the FCFS policies can be arbitrarily bad in moderate to heavy traffic conditions.

4.2 Partitioning Policies Based on Cyclic Queues

The intuitive reason the FCFS policies are unstable for high utilization values is that the average distance traveled per service, $\overline{d}$, remains fixed. The stability condition (16), however, implies $\overline{d} < \frac{1-\rho}{\lambda}$, so $\overline{d}$ must decrease as ρ (and λ) are increased. A policy that is stable for all values of ρ must, therefore, increasingly restrict the distance the server is willing to travel between services as the traffic intensity increases. This section examines two policies that achieve this restriction 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.

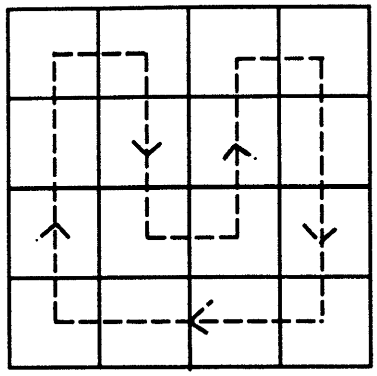

Figure 1: Sequence for Serving Subregions PART1 ($m = 4$)

4.2.1 The PART1 Policy

Consider the following policy for the DTRP which we call PART1: 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 Figure 1. (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 the 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 than the projection policy (e.g. moving directly to the first customer in the next subregion). The total travel distance around the subregions is $m^2(\sqrt{A}/m) = m\sqrt{A}$.

Note that to construct the pattern shown in Figure 1, m must be even. If m is odd, the

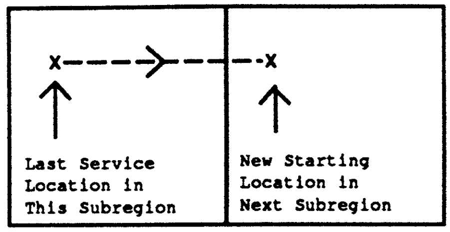

Figure 2: PART1 Projection Policy for Moving to Adjacent Subregion

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 $\overline{s^2} + 2c_1\frac{\sqrt{A}}{m} + c_2\frac{A}{m^2}$ 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 [4] still holds. Recalling that the expression in (3) is for the waiting time in queue only, the average system time for this policy is then given by

$$T_{PART1} = \frac{\lambda(\overline{s^2} + 2c_1\bar{s}\frac{\sqrt{A}}{m} + c_2\frac{A}{m^2})}{2(1 - \lambda(\bar{s} + c_1\frac{\sqrt{A}}{m}))} + \frac{1 - \frac{\lambda}{m^2}(\bar{s} + c_1\frac{\sqrt{A}}{m})}{2(1 - \lambda(\bar{s} + c_1\frac{\sqrt{A}}{m}))} m\sqrt{A} + c_1\frac{\sqrt{A}}{m} + \bar{s}. \qquad (23)$$

The stability condition for this policy is

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

If we define the critical value m_c by

$$m_c \equiv \frac{c_1\lambda\sqrt{A}}{1-\rho}, \qquad (24)$$

then the stability conditions says that m must be strictly larger than 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 this partitioning policy, we have the following theorem.

Theorem 2 *There exists a policy that has a finite waiting time for all $\rho < 1$ (the PART1 policy) for the DTRP and hence there exists an optimal policy for all $\rho < 1$.*

For given system parameters λ, $\overline{s}$, $\overline{s^2}$ and A, one could perform a one dimensional optimization over $m > 1$ using (23) to get the optimum number of partitions; however, since equation (23) is quite complicated, we concentrate on finding the optimal value of m for the light and heavy traffic cases. In the light traffic case, (23) becomes, approximately,

$$T_{PART1} \approx \frac{m\sqrt{A}}{2} + \frac{c_1\sqrt{A}}{m} + \overline{s}. \tag{25}$$

It is easy to check that the best value of $m > 1$ is $m = 2$ in this case. For $m = 2$, the waiting for service is approximately $1.26\sqrt{A} + \overline{s}$. This is more than either of the FCFS policies, so we gain nothing by using this policy in light traffic.

In heavy traffic ($\rho \rightarrow 1$), (24) implies that any feasible m is large ($m > m_c$); therefore ignoring the constant terms and assuming m is large, (23) becomes, approximately,

$$T_{PART1} \approx \frac{\lambda\overline{s^2} + m\sqrt{A}}{2(1-\rho-\lambda c_1\sqrt{A}/m)} = \frac{m^2\sqrt{A} + m\lambda\overline{s^2}}{2(m(1-\rho) - \lambda c_1\sqrt{A})}. \tag{26}$$

Differentiating the above with respect to m gives

$$\frac{dT_{PART1}}{dm} = \frac{\sqrt{A}}{2}\frac{m^2(1-\rho) - 2\lambda c_1\sqrt{A}m - \lambda^2 c_1\overline{s^2}}{(m(1-\rho) - \lambda c_1\sqrt{A})^2}.$$

Setting the numerator equal to zero and solving for m gives

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

Only the positive root is feasible, so for $\rho \rightarrow 1$ the second term under the radical approaches zero, and therefore

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

Thus, in heavy traffic, the optimal number of partitions approaches twice the critical number m_c.

If we substitute the optimal value m^* into (26), the optimal waiting time in heavy traffic is given by

$$T_{PART1} \approx 2c_1\frac{\lambda A}{(1-\rho)^2} + \frac{\lambda\overline{s^2}}{1-\rho}. \tag{27}$$

For $\rho \rightarrow 1$, the first term above dominates; therefore (recalling the bound (19)) we have

$$\frac{T_{PART1}}{T^*} \leq 8c_1, \qquad \text{as } \rho \rightarrow 1. \tag{28}$$

Since $c_1 \approx 0.52$, this says that the PART1 policy with $m = \lceil 2m_c \rceil$ has an average system time no more than approximately four times the optimal system time in heavy traffic. Also

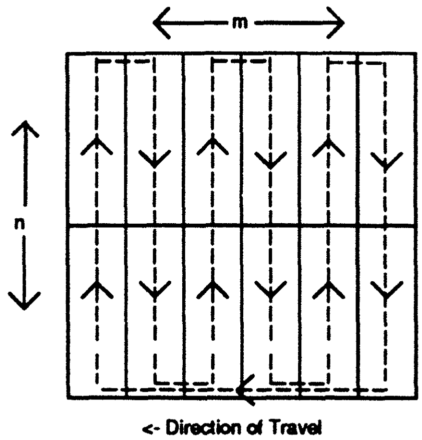

Figure 3: Sequence for Serving Subregions PART2 ($m = 6$, $n = 2$)

note that the dominating term is a constant times $\frac{\lambda A}{(1-\rho)^2}$, as suggested by the lower bound (19).

4.2.2 The PART2 policy

A modification of the PART1 policy, which we call PART2, has a performance guarantee in heavy traffic slightly better than $8c_1$. Consider the partition of the region $\mathcal{A}$ shown in Figure 3. The square has been divided horizontally into n equal sized strips and vertically into m equal sized strips to form nm, equal-sized rectangles. Within each subregion, the server travels on the vertical line centered in the rectangle and travels horizontally to service customers as shown in Figure 4. Only customers already in the subregion at the time of the server's arrival are serviced (gated service). Customers are served in the order in which they are encountered as the server moves along the vertical line. The sequence the server follows is indicated in Figure 3. It is easy to calculate that the total length of the path is $(m + 2 - \frac{2}{m})\sqrt{A}$. Again, we assume m is even in this calculation. For m odd, a slightly different pattern must be followed. Since we are primarily concerned with the heavy traffic case, m is large and the difference is negligible.

Note that the horizontal, roundtrip distances to individual customers are uniformly distributed with mean $\frac{\sqrt{A}}{2m}$ and second moment $\frac{A}{3m^2}$. If we consider these horizontal, roundtrip distances as the travel component of a customer's total service time, the system is *almost* a gated, cyclic queue. The problem is that the server's travel is distributed *between* the

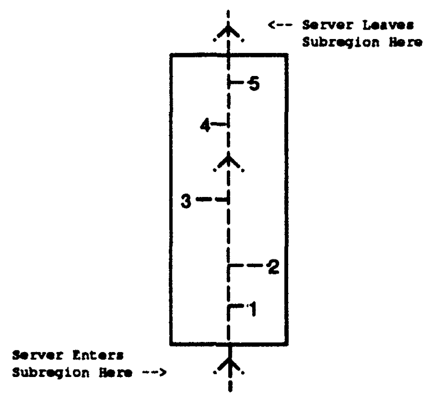

Figure 4: Policy for Serving Customers Within Subregion PART2 (Server travels horizontally to service customers. Numbers indicate service order.)

services within a subregion rather than occurring strictly before or after the services are completed; however, we can bound the delay as follows: suppose we consider a system where, upon arrival of the server in a subregion, the customers are "slid" towards the server in the vertical direction as shown in Figure 5. The server then services these customers (in the same order as before) prior to crossing the region. Let W_l denote the waiting time in queue for this modified policy and W_{PART2} denote the waiting time in queue for the PART2 policy. Clearly, the waiting time of customers cannot be lengthened by this procedure since we are serving each customer sooner than in the original system. Also, note that if we were to observe a busy period in a subregion, the total duration of the busy period as well as the set of customers served would be identical under both policies; therefore it follows that $W_l \leq W_{PART2}$.

Similarly, we can consider a system where the customers are "slid" away from the server (see Figure 5). The average waiting time in this system, denoted W_u is an upper bound on the waiting time in the original system, $W_u \geq W_{PART2}$. Now, each of these modified systems is in fact a cyclic queue. Further, it follows from the construction of the systems that for any realization of arrivals and on-site services, the *same* set of customers is served in each busy period in both system; therefore, since the customers in the upper bound system have to wait an additional travel time of exactly $\sqrt{A}/n$, $W_u = W_l + \frac{\sqrt{A}}{n}$. Thus,

$$W_l \leq W_{PART2} \leq W_l + \frac{\sqrt{A}}{n}. \tag{29}$$

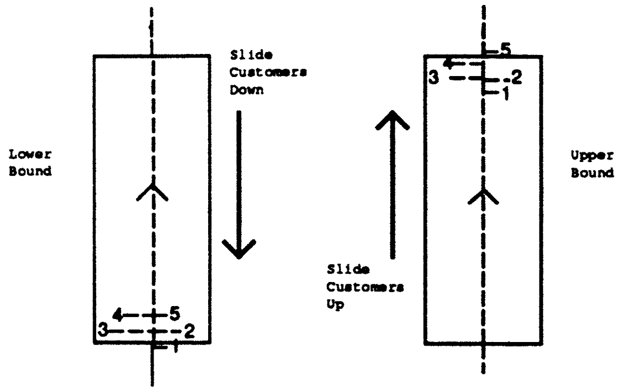

Figure 5: Lower and Upper Bound Systems for PART2 Policy

Since the lower bound system is a cyclic queue with gated service we can apply (4) to get

$$W_l = \frac{\lambda(\overline{s^2} + \overline{s}\frac{\sqrt{A}}{m} + \frac{A}{3m^2})}{2(1-\rho-\lambda\frac{\sqrt{A}}{2m})} + \frac{1+(\rho+\lambda\frac{\sqrt{A}}{2m})/nm}{2(1-\rho-\lambda\frac{\sqrt{A}}{2m})}(m+2-\frac{2}{m})\sqrt{A}. \tag{30}$$

If we let $n \to \infty$, (29) implies $W_{PART2} \to W_l$ and therefore from (30) we have

$$W_{PART2} = \frac{\lambda(\overline{s^2} + \overline{s}\frac{\sqrt{A}}{m} + \frac{A}{3m^2})}{2(1-\rho-\lambda\frac{\sqrt{A}}{2m})} + \frac{(m+2-\frac{2}{m})\sqrt{A}}{2(1-\rho-\lambda\frac{\sqrt{A}}{2m})}, \qquad n \to \infty. \tag{31}$$

Note that the lower and upper bounds in (29) are minimized for $n \to \infty$, so letting n get arbitrarily big is not only analytically convenient but also optimal. Intuitively, the behavior of the PART2 policy for large n is for the server to patrol the m vertical strips and to travel horizontally to service any customer encountered along the way.

As before, this policy can be shown to be no better than the FCFS policies in light traffic. In heavy traffic, we get results paralleling those for PART1. In order for the policy to be stable the critical number of vertical partitions is given by

$$m_c \equiv \frac{\lambda\sqrt{A}}{2(1-\rho)},$$

and in heavy traffic the optimal number of partitions can be shown to approach $2m_c$. Since $T_{PART2} = W_{PART2} + \overline{s} + \frac{\sqrt{A}}{2m}$, the average system time in heavy traffic becomes

$$T_{PART2} \approx \frac{\lambda A}{(1-\rho)^2} + \frac{\lambda\overline{s^2} + 2\sqrt{A}}{1-\rho},$$

where we have neglected constant terms . Again, since the first term dominates in heavy traffic we can use (19) to get

$$\frac{T_{PART2}}{T^*} \leq 4 \qquad \rho \rightarrow 1.$$

This is a slightly better performance guarantee compared with the PART1 policy (about 4% lower). Nevertheless, these performance guarantees are not terribly satisfactory from a practical standpoint. In the next section, we examine a policy that is quite different from the partitioning policies above and yields a significantly better performance guarantee.

4.3 The Traveling Salesman Policy

The travelling salesman policy (TSP for short) is based on collecting customers into sets that can then be served using an optimal TSP tour. Let N_k denote the kth set of N customers to arrive, where N is a given constant that parameterizes the policy, e.g. N_1 is the set of customers $1, \ldots, N$, 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 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 N_2 customers have arrived when the tour of N_1 is completed, they are serviced using a TSP tour; otherwise, the server waits until all N_2 customers arrive before serving it. In this manner, sets are serviced in a FCFS order. Note also that queueing of sets can occur.

Suppose one considers the set 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 $E[L_N] + N\overline{s}$ and the variance is $var(L_N) + N\sigma_s^2$, where $\sigma_s^2 = \overline{s^2} - \overline{s}^2$ 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

$$W_{set} \leq \frac{\frac{\lambda}{N}(\frac{N}{\lambda^2} + var[L_N] + N\sigma_s^2)}{2(1 - \frac{\lambda}{N}(E[L_N] + N\overline{s}))} \tag{32}$$

$$= \frac{\lambda(1/\lambda^2 + \frac{var[L_N]}{N} + \sigma_s^2)}{2(1-\rho-\lambda\frac{E[L_N]}{N})}. \quad (33)$$

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 (9) and (10)

$$\frac{E[L_N]}{N} \approx \beta_{TSP}\frac{\sqrt{A}}{\sqrt{N}} \quad (34)$$

and

$$\frac{var[L_N]}{N} \approx 0, \quad (35)$$

where the approximations are arbitrarily good 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

$$W_{set} \approx \frac{\lambda(1/\lambda^2 + \sigma_s^2)}{2(1-\rho-\lambda\beta_{TSP}\frac{\sqrt{A}}{\sqrt{N}})}. \quad (36)$$

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

$$N > \frac{\lambda^2\beta_{TSP}^2 A}{(1-\rho)^2}. \quad (37)$$

Thus in heavy traffic N *must* be large, so the asymptotic TSP results are indeed appropriate.

The waiting time given in (36) 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 (36) 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 prove that the average wait for a customer's set to form is $\frac{N-1}{2\lambda} \leq \frac{N}{2\lambda}$. By doing the same conditioning and noting that the travel time around the tour is no more that the length of the tour itself, it is easy to show that the expected wait for service once a customer's set enters service is no more than $\beta_{TSP}\sqrt{NA} + \frac{1}{N}\sum_{n=1}^{N} n\bar{s} \leq \beta_{TSP}\sqrt{NA} + \frac{N}{2}\bar{s}$. Therefore, the total system time, T_{TSP}, is

$$T_{TSP} \leq \frac{\lambda(1/\lambda^2 + \sigma_s^2)}{2(1-\rho-\lambda\beta_{TSP}\frac{\sqrt{A}}{\sqrt{N}})} + \frac{N(1+\rho)}{2\lambda} + \beta_{TSP}\sqrt{NA}. \quad (38)$$

We would like to minimize (38) with respect to N to get the least upper bound. (One can verify that (38) is convex, so there is indeed a minimum.) Doing this directly is tedious; therefore, motivated by the partitioning policy results, consider a change of variable where we pick N to satisfy

$$1-\rho-\lambda\beta_{TSP}\frac{\sqrt{A}}{\sqrt{N}}=\frac{1}{M}(1-\rho),$$

where $M>1$ is a new parameter. Such a value for N satisfies the stability condition and, in terms of M is given by

$$N=\frac{\lambda^2\beta_{TSP}^2A}{(1-\rho)^2(1-\frac{1}{M})^2}. \tag{39}$$

Thus, $M=2$ corresponds to N four times its critical value, and large values of M correspond to N close to its critical value. Substituting this value into (38) gives

$$T_{TSP}\leq\frac{\lambda\beta_{TSP}^2A(1+\rho)}{2(1-\rho)^2(1-\frac{1}{M})^2}+M\frac{\lambda(1/\lambda^2+\sigma_s^2)}{2(1-\rho)}+\frac{\lambda\beta_{TSP}^2A}{(1-\rho)(1-\frac{1}{M})}.$$

Letting $\rho\rightarrow 1$ the dominating term above is the first term, which approaches

$$\frac{\lambda\beta_{TSP}^2A}{(1-\rho)^2(1-\frac{1}{M})^2},$$

and therefore using the lower bound (19)

$$\frac{T_{TSP}}{T^*}\leq\frac{4\beta_{TSP}^2}{(1-\frac{1}{M})^2}\qquad\rho\rightarrow 1.$$

Thus, if we pick M to be large (which is equivalent to setting N close to its critical value (37)), the bound can be made arbitrarily close to $4\beta_{TSP}^2$ (e.g. select $M=(1-\rho)^{-1/2}$ and let $\rho\rightarrow 1$). Since the best estimate to date of β_{TSP} is approximately 0.72 [11], the TSP policy has a system time in heavy traffic of no more than approximately 2.1 times the optimal system time. This is an average system time about half that of the partitioning policies of the previous sections. It is encouraging that the (considerable) extra effort involved in constructing TSP tours does indeed yield a significant performance improvement over the less "intelligent" partitioning policies.

These results suggest that the policy of forming successive TSP tours, which is a reasonable policy in practice, is also quite good theoretically. Note that it is not necessary to have $\rho\approx 1$ for the $GI/G/1$ upper bound to be tight; what is really necessary is $N>>1$ in (37). For example, consider the case where $\overline{s}$ is small, and thus even moderate values of ρ result in very large arrival rates. In such cases, the optimal N is also close to the critical value; therefore, the queue of sets is truly in heavy traffic even though ρ is not close to one. This illustrates that preconceived ideas of what values of ρ constitute "light" and "heavy" traffic do not necessarily apply to the DTRP. A better guideline for applying the TSP policy, or any of the DTRP policies, is to use (19) to determine whether or not the minimum average number of customers in the system is large. If the minimum number is much larger than one, the partitioning, TSP or (as shown next) nearest neighbor policies are appropriate; if it is order one or smaller, the FCFS policies are probably more appropriate.

4.4 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 to obtain the heavy traffic lower bound (19), and 2) the shortest processing time (SPT) rule is known to be optimal for the classical $M/G/1$ queue [8]. 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 certainly not applicable; they are only suggestive.

Because of the dependencies among the travel distances d_i, analysis of the nearest neighbor policy is difficult. However, if we assume there exists a constant γ such that

$$E[d_i|N] \leq \gamma \frac{\sqrt{A}}{\sqrt{N}}, \tag{40}$$

where N is the number of customers in the system at a service completion epoch, then by using a modification of the argument in [14] Section 5.5, it is possible to show that for $\rho \rightarrow 1$,

$$T_{NN} \leq \frac{\gamma^2 \lambda A}{(1-\rho)^2}. \tag{41}$$

where T_{NN} denotes the system time of the nearest neighbor policy. As a result,

$$\frac{T_{NN}}{T^*} \leq 4\gamma^2 \qquad as \quad \rho \rightarrow 1.$$

Though we believe that assumption (40) is quite reasonable, there is no analytical guidance in selecting the constant γ. In the lower bound (19) we made an assumption similar to (40); however, in that case we argued that the constant 1/2 obtained from assuming uniformly distributed points gave a lower bound. In the case of the nearest neighbor policy, no analogous argument exists for a constant γ that gives an upper bound.

We therefore performed simulation experiments to estimate γ and to verify the asymptotic behavior of T_{NN}. The method of ***batch means*** (see [17]) was used to estimate the

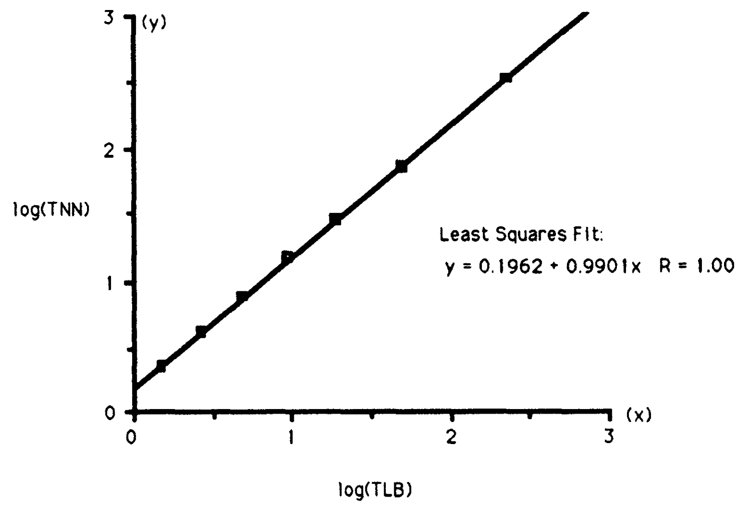

Figure 6: Simulation Estimate of T_{NN} vs. Lower Bound (log-log plot)

graphed separately.

Figure 7 shows the system time as a function of ρ for the FCFS and SQM policies. The lower bound is also included. Note that although the SQM policy is asymptotically optimal as $\rho \rightarrow 0$, 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 system times for the PART1, PART2, TSP and NN policies were calculated (simulated in the case of NN) for the same example. The results are shown in Figure 8 along with the lower bound (19). Note that the graphs have nearly identical shape as one would expect from the $\frac{\lambda A}{(1-\rho)^2}$ asymptotic behavior of each policy. (Only the constant of proportionality differs.)

The relative performance of the policies for intermediate values of ρ did not reflect the asymptotic performance in some cases. For example, in the range $0.1 < \rho < 0.8$, the PART1 policy actually performed better than the PART2 policy. Also, because we assumed a large set size for the TSP analysis, the upper bound for the TSP was plausible only for values of $\rho > 0.5$. The results suggested a mixed strategy for this example: for $\rho < 0.05$ the SQM policy was clearly the best. For $0.05 < \rho < 0.1$, the FCFS and NN performed comparably.

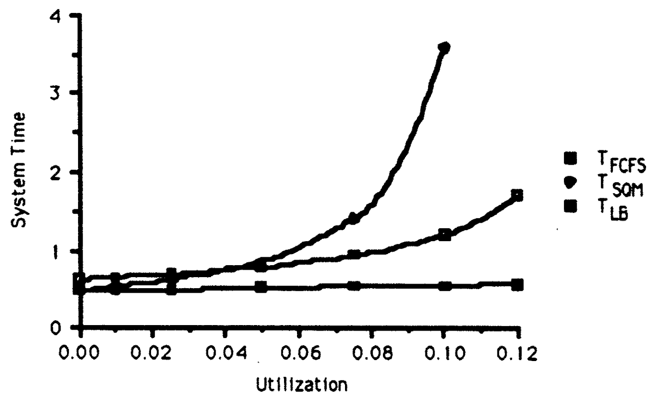

Figure 7: System Times for Light Traffic Policies: Numerical Example

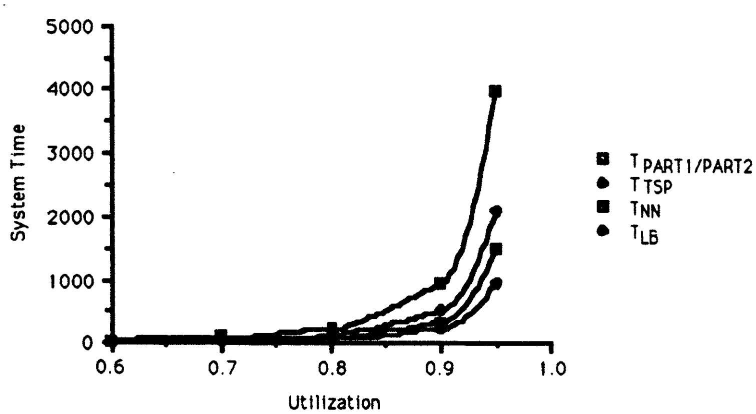

Figure 8: System Time for Heavy Traffic Policies: Numerical Example

As ρ was increased above 0.1, the NN policy dominated. The TSP was clearly second best in the range where its upper bound was valid ($\rho > 0.5$). This example suggests that no single policy dominates for all traffic conditions, and thus a case-by-case analysis is warranted.

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. In heavy traffic we showed that partitioning policies behave reasonably well, since they have constant factor performance guarantees, and have finite system times for all values of $\rho < 1$. In addition, they have the advantage of being easily extendible to the case of many servers (vehicles). The best policies in heavy traffic were the TSP and nearest neighbor (NN) policies, which were within a small constant factor of the lower bound. The TSP policy has the advantage that the server regularly returns to the depot. It also appears more "fair" in the sense that it partially obeys a FCFS discipline, since sets are served in FCFS order. It also has a provable performance guarantee. The nearest neighbor policy, on the other hand, performs about 25% better than the TSP strategy according to our simulation study; however, it can severely violate the FCFS discipline and does not return to a single location on a regular basis. Also, it does not have provable performance guarantees.

A common characteristic of all the policies we proposed is that they are easily implementable in a real-world environment. They also have, despite their diversity, identical asymptotic behavior in heavy traffic. The behavior 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. Its root cause is perhaps most clearly seen in the derivation of the bound (19). Though the proof of this bound is partly heuristic, the insights and results it provides are, in our estimation, correct.

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 (vehicle) models. This is a topic we are currently investigating. Another, probably more elusive, goal would be a rigorous proof of the bound on the nearest neighbor distance (18), which was the critical relation in establishing (19). Finally, one could certainly construct other DTRP policies and analyze them using the techniques of Section 1.

Acknowledgement

The first author is indebted to his colleague and friend Professor Harilaos Psaraftis for bringing dynamic vehicle routing problems to his attention and for interesting discussions.

References

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

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

[3] 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.

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

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

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

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

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

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

[10] 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.

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

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

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

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

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

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

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

[18] 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.

[19] 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.

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

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

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

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

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