Stochastic and Dynamic Vehicle Routing in the Euclidean Plane: The Multiple-Server, Capacitated Vehicle Case

by

D. J. Bertsimas and G. van Ryzin

OR 224-90 August 1990

Stochastic and Dynamic Vehicle Routing in the Euclidean Plane: The Multiple-Server, Capacitated Vehicle Case

Dimitris J. Bertsimas * Garrett van Ryzin †‡

August 1990

Abstract

In a previous paper [12], we introduced a new model for stochastic and dynamic vehicle routing called the dynamic traveling repairman problem (DTRP), in which a vehicle traveling at constant velocity in a Euclidean region must service demands whose time of arrival, location and on-site service are stochastic. The objective is to find a policy to service demands over an infinite horizon that minimizes the expected system time (wait plus service) of the demands. We showed that the stability condition did not depend on the geometry of the service region (i.e. size, shape, etc.). In addition, we established bounds on the optimal system time and proposed an optimal policy in light traffic and several policies that have system times within a constant factor of the lower bounds in heavy traffic. We showed that the leading behavior of the optimal system time had a particularly simple form which increases much more rapidly with traffic intensity than the system time in traditional queues (e.g. M/G/1).

In this paper, we extend these results in several directions. First, we propose new bounds and policies for the problem of m identical vehicles with unlimited capacity and show that in heavy traffic the system time is reduced by a factor of $1/m^2$ over the single server case. Policies based on dividing the service region into m equal subregions

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

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

‡The research of both authors was partially supported by the National Science Foundation under grant DDM-9014751.

served independently by one vehicle each are shown to have the same constant factor guarantees as in the one vehicle case. We then consider the case in which each vehicle can serve at most q customers before returning to a depot. In this case, the stability condition depends strongly on the geometry of the region. Several policies that have system times within a constant factor of the optimum in heavy traffic are then proposed. We also extend these results to higher dimensions and general locational distributions. Finally, we propose and analyze several optimization problems (e.g. fleet sizing and districting) that arise in the context of dynamic vehicle routing.

Key words. dynamic vehicle routing, multiple servers, capacitated vehicles, traveling repairman problem, traveling salesman problem, queueing, bounds, heuristics.

1 Introduction

The classical approach to vehicle routing problems is to view them as static and deterministic. A set of known customer locations defines an instance, and the objective is to visit customers so as to minimize the total travel cost subject to certain constraints (e.g. a limit on vehicle capacity). These classical problems have generated significant research interest over the years (see for example [24], [16]) resulting in major contributions in the areas of combinatorial optimization, the analysis of heuristics and complexity theory. However, as models for the type of vehicle routing problems encountered in practice they are not always appropriate. Many real-life problems involve considerable uncertainty in the problems data. For example, locations may be known only probabilisticaly in advance and the demands they place on vehicle capacity may be random. In addition, requests for service often arrive sequentially in time, and again these arrivals epochs may be stochastic. Finally, the objective of minimizing travel distance is not necessarily paramount; in a dynamic setting, the delivery time (wait for service) is often a more appropriate objective.

As a canonical example of an application with these characteristics, consider the following utility repair problem: A utility firm (electric, gas, water and sewer, highway, etc.) is responsible for the maintenance of a large, geographically dispersed facilities network. The network is subject to failures which occur randomly both in time and space (location). The firm operates a fleet of repair vehicles which are dispatched from a depot to respond to failures. The vehicle crews spend a random amount of time servicing each failure before they are free to move on to the next failure. The firm would like to operate its fleet in a way that minimizes the average downtime due to failures.

There are many closely related problems to this canonical example that arise in practice. For example, consider a firm that delivers a product from a central depot to customers based on orders that arrive in real-time. Orders are queued and delivered by a fleet of vehicles with the objective of minimizing the average wait for delivery. In still other applications such as emergency service system operations (e.g. fire, ambulance and police) or courier services similar characteristics and criteria apply.

1.1 Literature Survey

The complexities of such problems are often incorporated in the classical framework through the use of rolling horizon procedures. These involve planning routes for a fixed period into the future, often with the option of adding or deleting demands and modifying routes as time advances. See Brown and Graves [13], Powell [27] and Psaraftis [29] for examples of this approach. Though useful for data–intensive tactical problems, they are inherently *ad hoc* and do not give the insight necessary for strategic planning.

Several researchers have proposed alterative models that explicitly consider some combination of stochastic, dynamic demands or congestion/waiting time measures. A static, deterministic problem which uses the waiting time objective is the traveling repairman problem (TRP). In the TRP a vehicle services a set of n demands starting from a depot. The distances between demands i and j, $d(i,j)$ are given so if the sequence followed is $(1,2,\ldots,n,1)$ the total waiting time is $\sum_{i=1}^{n} w_i$ where w_i is the waiting time of demand i given by $w_i = \sum_{j=1}^{i-1} d(j,j+1)$. The problem is to minimize the total waiting time. The problem is known to be NP-hard (see Shani and Gonzalez [31] and Afrati *et al.* [1]). The TRP even seems difficult on trees. Minieka [25] proposes an exponential $O(n^p)$ for the TRP on a tree $T=(V,E)$ where $|V|=n$ and p is the number of leaves. Unfortunately, little else is known about the problem.

Jaillet [19], Bertsimas [8], [9] and Bertsimas, Jaillet and Odoni [10] address uncertainty in demand locations in their formulation of the probabilistic traveling salesman problem (PTSP) and the probabilistic vehicle routing problem (PVRP). In the PTSP there are n given points, and on any given instance of the problem only a subset $\mathcal{S}$ consisting of $|\mathcal{S}|=k$ of the points must be visited. Given the probability of each instance $p(\mathcal{S})$, we wish to find *a priori* a tour through all n points, where on a given instance the k points will be visited in the order of this tour. The problem is to find such a tour that is of minimum length in the expected value sense. In the case where the vehicle has capacity q, the resulting problem is the PVRP. Though the problem is stochastic and can model the real-time occurrence of problem instances, the strategy is inherently static and is solved using only probabilistic information.

Dynamic and stochastic characteristics have been considered in the context of location problems by Batta *et al.* [5] and Berman *et al.* [7] who define the stochastic queue median problem (SQMP). In the SQMP demands arrive to nodes of a network according to independent Poisson processes. The demands require a generally distributed amount of service from a vehicle based at a depot that follows a first-come-first-serve (FCFS) order, returning to the depot after each service completion. Thus, the system operates as an M/G/1 queue with a service time distribution that depends on the depot location. The problem is to locate the depot so as to minimize the expected waiting time. The model is well suited to emergency service applications (e.g. police, fire and ambulance dispatching), but the service strategy is quite restrictive and less appropriate for delivery and repair problems.

A somewhat closer representation of our canonical example is found in the polling system and machine repairman literature. A polling system in defined identically to the SQMP except that the service strategy is to repeatedly visit nodes according to a fixed permutation. The server either serves all customers present at a node at the time of arrival (gated service) or serves a node until no customers are left (exhaustive service) before moving to the next node in the sequence. (See Takagi [34] for a comprehensive survey of polling systems.) The policy can be enriched by using a general polling table where sequences longer and more complicated than simple, cyclic permutations are used [3]. Unfortunately, even the performance analysis of such systems is quite difficult and often involves solving large systems of linear equations [15]. Browne and Yechiali [14] obtain dynamic index rules for visiting nodes based on optimization over a limited horizon of one cycle; however, the approach requires the distances to be decomposable so that $d(i,j) = d_i + d_j$ for all i and j, which maybe appropriate for some computer system and manufacturing applications but is unrealistic in a vehicle routing context. In addition, it is not clear how their myopic criterion relates to the objective of minimizing average waiting time over an infinite horizon.

The machine repairman problem, a closely related problem, has the same network structure as in the SQMP and polling systems, but the capacity at each node is one. Thus individual nodes can be thought of as single machines that fail randomly in time and wait to be repaired by a traveling repairman. (See Stecke and Aronson [33] for a review.) Agnihothri [2] solves only the perfectly symmetric case (i.e. identical node statistics and identical travel times between all nodes) exactly using a Semi-Markov model. Due to the symmetry, however, all work conserving policies are equally good, so this formulation fundamentally avoids the issue of optimization. In addition, the resulting performance measures are quite complex. As a result, little insight is gained for the realistic asymmetric case.

A formulation that closely matches our canonical application is the dynamic traveling salesman problem (DTSP) proposed by Psaraftis [28]. As in the other network formulations, he considers a complete graph with n nodes each of which receives arrivals according to independent Poisson process. There is a general service time distribution for each node. The distance between nodes is known. The arrivals are serviced by a vehicle traveling on the network, and the goal is to optimize over some performance measure such as throughput or waiting time. This model motivated our initial investigation; however, it seems inherently complex, and general results have yet to be obtained.

1.2 Review of Single Server DTRP Results and Policies

In [12] we proposed and analyzed a new model for dynamic and stochastic vehicle routing problems. The model, called the dynamic traveling repairman problem (DTRP), incorporates essentially all the features of the canonical network repair application mentioned above. It is defined as follows: demands for service arrive according to a Poisson process with rate λ to a Euclidean service region $\mathcal{A}$ of area A. We assume only that the region $\mathcal{A}$ is a bounded, convex set with piecewise linear boundaries. Upon arrival, demands are assigned an independent and uniformly distributed location in $\mathcal{A}$. Demands are serviced by a vehicle that travels at constant velocity v, and each demand requires a generally distributed, on-site service time with finite first and second moments denoted by $\overline{s}$ and $\overline{s^2}$ respectively. The objective is to find policies for routing the vehicle that minimize the average system time (waiting time plus on-site service time) of demands.

Using results from geometrical probability, queueing theory and combinatorial optimization, we were able to obtain several interesting results for this problem. It is useful to review the main results at this point as they will be used repeatedly in the remainder of the paper.

First, in the light traffic case ($\lambda \rightarrow 0$),we showed the following policy is optimal:

- **The Stochastic Queue Median (SQM) Policy:** Locate the server at the median, x^*, of $\mathcal{A}$ and serve customers in FCFS order, returning to the median after each service.

In this case we showed that the expected optimal system time, T^* satisfies

$$T^* \sim \frac{E[\|X - x^*\|]}{v} + \overline{s} \qquad \text{as } \lambda \rightarrow 0,$$

where X is a uniform location in $\mathcal{A}$. (Note that the first term above is simply the expected travel time to the median). However, this policy quickly becomes unstable as the traffic intensity increases.

In heavy traffic, we discovered a quite different and unexpected behavior. If we let $\rho \equiv \lambda \overline{s}$ denote the fraction of time the vehicle spends in on-site service, then we showed that policies exist that have finite system times for all $\rho < 1$. This is surprising in that the condition is completely independent of the service region size and shape. It is also the mildest stability restriction one could hope for.

We then showed that there exists a constant γ such that

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

(We extend this result to the m vehicle case in §2 below.) Note that this grows like $(1-\rho)^{-2}$. Thus, though the stability condition is similar to a traditional queue, the system time increases more rapidly as $\rho \to 1$.

We constructed several policies that have finite system times for all $\rho < 1$. In addition, we showed that these policies have the same asymptotic behavior, namely

$$T \sim \gamma_\mu^2 \frac{\lambda A}{v^2(1-\rho)^2} \qquad \text{as } \rho \to 1,$$

where the constant γ_μ demands only on the policy μ. Hence, by comparing this to the lower bound above, we see that the policies have a constant factor performance guarantee in heavy traffic relative to the optimum value T^*. The policies and their associated constants are given below. When exact analysis failed or only bounds were possible, empirical values for the constants were obtained via simulation.

- **The Partition (PART) Policy:** (This policy applies to the case where $\mathcal{A}$ is a square.) Divide the region into n equal subregions which are served sequentially such that each subregion is adjacent to the next subregion in the sequence except, perhaps, for the last one. (See [12] for details.) Within each subregion, service demands in FCFS order until no more demands are left (exhaustive service). Then move on to the next subregion in the sequence. Optimize over n according to [12].

  Constant Value: $\gamma_{PART} = \sqrt{2c_1} \approx 1.02$,

  (c_1 is the expected distance between two uniform points in the unit square.)

- **The Traveling Salesman (TSP) Policy:** As demands arrive, form them into sets of size n. When all n demands have arrived, consider it the arrival of a set. Service sets in FCFS order by forming a TSP tour on the set of demands. Optimize over n according to [12].

Constant Value: $\gamma_{TSP} = \beta_{TSP} \approx 0.72$,

(β_{TSP} is the asymptotic TSP constant as in [6].)

- **The Space Filling Curve (SFC) Policy:** A space filling curve is a mapping of locations from the unit square to positions (*preimages*) on the interval [0,1] (see [4] and [26]). Using a real-time sorted list, maintain the preimages of all outstanding demands in the system. Visit demands according to the order in which they are encountered in continuous sweeps of the interval [0,1].

  Provable Bound on Constant Value: $\gamma_{SFC} \leq 2$,

  Empirical Value: $\gamma_{SFC} \approx 0.66$.

- **The Nearest Neighbor (NN) Policy** After each service completion, serve next the demand that is closest to the vehicle.

  Empirical Value: $\gamma_{NN} \approx 0.64$.

It is remarkable that this diverse collection of policies have identical asymptotic behavior. Moreover, this behavior has a particularly simple form that clearly shows how the average system time behaves as a function of vehicle velocity, service region size, on-site service statistics and traffic intensity. As mentioned, the behavior is distinctly different from that found in traditional queueing systems. In particular, it is proportional to $(1-\rho)^{-2}$ rather that $(1-\rho)^{-1}$ and also does not depend on the on-site service time variability (at least in its leading behavior).

1.3 Overview and Contributions of this Paper

As satisfying as these results are, the model of a single uncapacitated vehicle is somewhat unrealistic for most practical purposes. Therefore, we were motivated to expand the analysis to more realistic configurations. In the remainder of the paper, we show how the results above can ultimately be extended to the case where the region $\mathcal{A}$ is serviced by a homogeneous fleet of m vehicles operating out of a set $\mathcal{D}$ of $|\mathcal{D}| = m$ depots, where each vehicle is restricted to visiting at most q customers before returning to its respective depot. (The depot locations need not be distinct.) We show that the minimum expected system time, T^*, in this case has the following lower bound:

$$T^* \geq \gamma'^2 \frac{\lambda A(1 - \frac{1}{q^2})}{m^2 v^2 (1 - \rho - \frac{2\lambda \bar{r}}{mqv})^2} - \frac{m(1-2\rho)}{2\lambda},$$

where $\rho \equiv \lambda\bar{s}/m$, $\bar{r}$ denotes the expected distance from a uniform location in $\mathcal{A}$ to the closest point in $\mathcal{D}$ and γ' is a numerical constant. Note that for the case $q \to \infty$ and $m = 1$, this reduces to our earlier result.

When $m > 1$ and $q \to \infty$, we show that policies with the same constant factor performance guarantee as in the single server case can be constructed by simply partitioning $\mathcal{A}$ into m equal subregions and serving each one independently using one of the single server policies mentioned above.

When q is finite, the above expression provides some intuitively satisfying insights. For example, consider the case where $m = 1$ and $q < \infty$. Then the above bound implies a stability condition of

$$\rho + \frac{2\lambda\bar{r}}{vq} < 1.$$

We show this condition is also sufficient by constructing policies with this same stability criterion and asymptotic behavior. Observe that this condition is no longer independent of the service region geometry in this case because of the presence of $\bar{r}$; however, for $q \to \infty$ the dependence vanishes.

The second term in the stability condition has the interpretation of a *radial collection cost* in the sense of Haimovich and Rinnoy Kan [18]. That is, $2\bar{r}/v$ is essentially the average time required to reach a set of q customers from the nearest depot (the radial cost). Dividing by q gives the average radial cost per customer, and hence multiplying by λ we obtain the fraction of time the server spends in radial travel. The above condition says that as long as this fraction plus the fraction of time spent on-site is less than one, the system will be stable. Furthermore, the waiting time grows like the inverse square of this stability difference just as it does in the uncapacitated case. Thus, we see the crucial role the average radial distance $\bar{r}$ plays in the system's stability. Indeed, we prove that if one has the option of locating the depot anywhere within $\mathcal{A}$, then minimizing $\bar{r}$ (i.e. locating the depot at the median) is always optimal in heavy traffic.

For $m > 1$ and q finite, we construct policies that have system times within a constant factor of the optimum for several cases. In the case where all m vehicles are based out of the same depot, we show that a policy based on subdividing the region into squares, forming tours of q customers within each square and then serving tours in FCFS order has a constant factor guarantee. When there are k depots, this policy can be applied provided certain symmetry conditions hold. Specifically, the k Voronoi cells must be identical and there must be $m = kp$ vehicles for some integer p. Also, when k is large and the depots occupy the k-median locations, then the Voronoi cells are approximately identical and similar policies can again be constructed.

For completeness, we also discuss extensions to regions in higher dimensions and to locational distributions that are not uniform. In addition, we use our results to formulate and analyze several strategic planning problems such as optimization of travel and waiting cost, optimal fleet sizing and optimal districting.

The remainder of the paper is organized as follows: In §2 we examine the case of m vehicles without capacity. We first obtain lower bounds through an extension of our earlier results [12]. This serves to establish several key lemmas necessary for the q capacitated case. We then propose several policies for this case and analyse their performance with respect to the lower bounds. In §3 we examine the capacitated case. We establish a lower bound and propose and analyze a policy for the single depot, multi-vehicle case. As mentioned, this policy can be extended to the multi-depot case provided certain symmetry conditions hold and these are discussed in §3 as well. In §4 we extend our results to regions in higher dimensions and to nonuniformly distributed demands. In §5 we introduce and analyze several strategic planning models based on the results of §1 and §2. Finally, in §6 we give our conclusions.

2 The m-Vehicle, ∞-Capacity DTRP

2.1 Lower Bounds

Before deriving the lower bounds, some notation and more precise definitions are needed. We shall index demands according to their *service* order. We let s_i denote the on-site service time of the i-th demand served, W_i denote the i-th demand's waiting time and $T_i = W_i + s_i$. With each demand we associate a travel distance, d_i, which is the distance the server travels in going from demand $(i-1)$ to demand i. The limiting expected values of these random variables are defined by $\bar{s} = \lim_{i\to\infty} E[s_i]$, $W = \lim_{i\to\infty} E[W_i]$, $T = \lim_{i\to\infty} E[T_i]$ and $\bar{d} = \lim_{i\to\infty} E[d_i]$. We shall assume these limits exist. Finally, as before we write T_μ to indicate the system time of a particular policy μ and T^* to indicate the optimal system time.

2.1.1 A Light Traffic Lower Bound

The first bound for this case is established by dividing the system time of demand i, T_i, into three components: the waiting time of demand i due to the servers travel prior to serving i, denoted W_i^d; the waiting time of demand i due to on-site service times of demands served prior to i, denoted W_i^s; and demands $i's$ own on-site service time, s_i. Thus,

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

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

$$T = W^d + W^s + \overline{s}, \tag{1}$$

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 W_i^d is at least the travel delay between the location of the closest server at the time of arrival and demand i's location. In general, the servers are located in the region according to some generally unknown spatial distribution that depends on the policy. However, suppose we had the option of locating the m servers in the best *a priori* set of location, $\mathcal{D}^*$; that is, the set of location that minimizes the expected distance to the demand's uniformly distributed location, X. This certainly yields a lower bound on the expected distance between the nearest server and the demand's location, and since this in turn is a lower bound on W^d we obtain

$$W^d \geq \frac{1}{v} \min_{|\mathcal{D}|=m} E[\min_{x_0 \in \mathcal{D}} ||X - x_0||]. \tag{2}$$

The set of locations that achieves the minimization above is called the set of *m-median* locations of the region $\mathcal{A}$.

Using the trivial bound $W_s \geq 0$ and combining with (1) and (2) we easily get the following theorem:

Theorem 1

$$T^* \geq \frac{1}{v} E[\min_{x_0 \in \mathcal{D}^*} ||X - x_0||] + \overline{s}.$$

This bound is most useful in the case of light traffic ($\lambda \to 0$).

2.1.2 A Heavy Traffic Lower Bound

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

Theorem 2 *There exists a constant γ such that*

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

<u>Proof</u>

We begin with an important lemma which will be used again in the q-capacitated case.

Lemma 1

$$\overline{d} \geq \gamma \frac{\sqrt{A}}{\sqrt{N + m/2}}, \tag{3}$$

where $\gamma = \frac{2}{3\sqrt{6\pi}}$ and N is the average number of customers in queue.

Before proving this lemma, observe that Theorem 2 is easily derived from it by substitution into the stability condition,

$$\overline{s} + \frac{\overline{d}}{v} \leq \frac{m}{\lambda}, \tag{4}$$

which yields

$$\overline{s} + \frac{\gamma\sqrt{A}}{v\sqrt{N + m/2}} \leq \frac{m}{\lambda}.$$

After rearranging, noting that $T = W + \overline{s}$ and $N = \lambda W$, we obtain the bound of Theorem 2. Thus, Theorem 2 is established once Lemma 1 is proven.

□ (Theorem 2)

<u>Proof of Lemma 1</u>

Consider a random "tagged" demand and define,

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

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

$X_0 \equiv$ *The tagged demand's location.*

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

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

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

$$P\{Z_i \leq z\} \leq \frac{\pi z^2}{A}, \tag{5}$$

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

The set of locations from which a server can visit the tagged demand is at most $\mathcal{S}_0 \cup \mathcal{S}_1$; therefore, the value of d_i for the tagged demand is at least $\min\{Z_0^*, Z_1, ..., Z_{N_1}\}$. Hence,

$$\overline{d} \geq E[\min\{Z_0^*, Z_1, ..., Z_{N_1}\}]. \tag{6}$$

We next bound the right hand side of (6). 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 (5), we obtain

$$E[\min\{Z_0^*, Z_1, ..., Z_{N_1}\}] \geq z(1 - P\{Z_0^* \leq z\} - N\frac{\pi z^2}{A}). \tag{7}$$

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

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

Proof

The proof relies on the following result due to Haimovich and Magnanti [17] 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 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 (c.f. [30]) is that Z^* is stochastically larger than 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 $E[N_0] = N + m$ establishes the lemma.

□ (Lemma 2)

Using the result of Lemma 2 in (7) yields

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

Maximizing the right hand side with respect to z gives

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

This establishes Lemma 1 with $\gamma = \frac{2}{3\sqrt{6\pi}} \approx 0.153$.

□ (Lemma 1)

2.2 An Optimal Light Traffic Policy

A direct extension of the SQM policy to the m-server case gives an optimal policy in light traffic as we now demonstrate. Consider the following policy:

- **The m Stochastic Queue Median (mSQM) Policy**: Locate one server at each of the m median locations for the region $\mathcal{A}$. When demands arrive, assign them to the nearest median location and its corresponding server. Have each server service its respective demands in FCFS order returning to its median after each service is completed.

Let $j = 1, \cdots, m$ index the m Voronoi cells, $\mathcal{A}_j$ denote the j-th cell, $A_j = |\mathcal{A}_j|$ and x_j^* denote the j-th median location. Also, let $\lambda_j = \frac{A_j}{A}\lambda$ denote the arrival rate to cell j and $\rho_j = \lambda_j \overline{s}$ server j's utilization. Finally, for a uniformly distributed location $X \in \mathcal{A}$ let

$$\overline{d_j} = E[||X - x_j^*|| \mid X \in \mathcal{A}_j]$$

and

$$\overline{d_j^2} = E[||X - x_j^*||^2 \mid X \in \mathcal{A}_j].$$

Note that each cell j is an independent, single-server SQM system operating as an $M/G/1$ queue with first moment $\overline{s} + 2\overline{d_j}/v$ and second moment $\overline{s^2} + 4\overline{s}\overline{d_j}/v + 4\overline{d_j^2}/v^2$. Since, the probability of a given arrival lands in cell j is simply A_j/A, we have that

$$T_{mSQM} = \sum_{j=1}^{m} \frac{A_j}{A} \frac{\lambda_j(\overline{s^2} + 4\overline{s}\overline{d_j}/v + 4\overline{d_j^2}/v^2)}{2(1 - 2\lambda_j \overline{d_j}/v - \rho_j)} + \sum_{j=1}^{m} \frac{A_j}{A}(\overline{d_j}/v + \overline{s}),$$

where the terms in the second sum are the weighted one-way travel time plus on-site service time means in each cell. As $\lambda \to 0$, $\lambda_j \to 0$ for all j and thus the contribution of the first term tends to zero, while the second term is simply $\frac{1}{v}E[\min_{x_0 \in \mathcal{D}^*} ||X - x_0||] + \overline{s}$ by construction since $\{x_j^*\} = \mathcal{D}^*$ and $\mathcal{A}_j = \{x | j = \text{argmin}_k ||x - x_k^*||\}$. Therefore,

$$T_{mSQM} \sim \frac{E[\min_{x_0 \in \mathcal{D}^*} ||X - x_0||]}{v} + \overline{s} \qquad \text{as } \lambda \to 0.$$

Hence, by comparing this to Theorem 1 we see the mSQM policy is asymptotically optimal in light traffic. One can verify from the individual stability conditions for each cell that if $A > 0$ there is a critical value $\rho_c < 1$ such that the system time is unbounded for $\rho \geq \rho_c$; therefore, in light of Theorem 2, it is clear that the mSQM policy has an unbounded cost relative to optimum for $\rho \to \rho_c$ and certainly for $\rho \to 1$.

2.3 Heavy Traffic Policies

We next examine several heavy traffic policies. We prove that policies based on randomized assignment of arrivals to servers are within a constant factor of the lower bound for $\rho \to 1$

for a given number of servers m but this factor increases with m. We show a similar result for the case of a $G/G/m$ version of the TSP policy. Finally, we show that if the service region $\mathcal{A}$ is divided into equal sized subregions and one of the single server policies is applied in each region, then the same constant factor over the optimum as in the single-vehicle case is achieved for any m. This leads to a conjecture that the optimal m server policy in heavy traffic is to apply the optimal single server policy to m equal sized subregions.

2.3.1 Randomized Assignment (RA)

One possible strategy is to split the Poisson input process through randomization into m sub-processes, one for each server, and then have servers individually service their stream using a heavy traffic, single-server policy μ. It is well known that if we assign each arrival independently to server j with probability $1/m$ for all j, then the resulting sub-streams will be independent Poisson processes each with rate λ/m. Therefore, each individual server sees a demand arrival process with rate λ/m and operates independently in the entire region $\mathcal{A}$ to service it. Since an arrival is equally likely to be assigned to any one of the m servers and each operates identically, the system time for randomized assignment is simply

$$T_{RA\mu} \sim \gamma_\mu^2 \frac{\lambda A}{mv^2(1-\rho)^2} \qquad \text{as } \rho \to 1.$$

where as before $\rho \equiv \lambda\bar{s}/m$. Comparing this to the bound in Theorem 2 implies

$$\frac{T_{RA\mu}}{T^*} \leq m\frac{\gamma_\mu^2}{\gamma^2}. \quad \text{as } \rho \to 1.$$

Thus, the performance guarantee for $RA\mu$ has the undesirable characteristic of increasing with the number of servers.

2.3.2 A G/G/m Version of the TSP Policy

One might expect that a more intelligent allocation of customers to servers might yield a better bound. Such is indeed the case as shown by the following G/G/m version of the TSP policy. We present the details of this analysis not only for the m server case but also to establish some analysis techniques and relations that will be useful when we discuss the q capacitated case and the optimization problems.

The TSP policy is based on collecting customers into sets that can then be served using an optimal TSP tour. The single server version operates as follows: 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. Note that if one considers sets as customers, this policy forms a G/G/1 queue. The interarrival distribution is Erlang of order n, and thus 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 $E[L_n]/v + n\overline{s}$ and the variance is $var(L_n)/v^2 + n\sigma_s^2$, where $\sigma_s^2 = \overline{s^2} - \overline{s}^2$ is the variance of the on-site service time.

In the G/G/m policy, we form sets in exactly the same manner only the sets are served in FCFS order by m vehicles so as to form a G/G/m queue. We next make use of a heavy traffic limit due originally to Kingman [21] for the waiting time W in a G/G/m queue (c.f. [23]),

$$W \sim \frac{\lambda(\sigma_a^2 + \sigma_b^2/m^2)}{2(1-\rho)} \qquad \text{as } \rho \to 1 \tag{8}$$

where σ_a^2 and σ_b^2 are the variances for the interarrival times and service time respectively, and $\rho = \lambda\overline{s}/m$. Letting W_{set} denote the waiting time of a set, this limit in our case gives

$$W_{set} \sim \frac{\frac{\lambda}{n}(\frac{n}{\lambda^2} + \frac{1}{m^2}(var[L_n]/v^2 + n\sigma_s^2))}{2(1 - \frac{\lambda}{mn}(E[L_n]/v + n\overline{s}))} \tag{9}$$

$$= \frac{\lambda(1/\lambda^2 + \frac{1}{m^2}(\frac{var[L_n]}{nv} + \sigma_s^2))}{2(1 - \rho - \frac{\lambda}{m}\frac{E[L_n]}{nv})}. \tag{10}$$

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 can apply the following asymptotic TSP results (c.f. [24], [32]):

$$\frac{E[L_n]}{n} \approx \beta_{TSP}\frac{\sqrt{A}}{\sqrt{n}}, \tag{11}$$

and

$$\frac{var[L_n]}{n} \approx 0, \tag{12}$$

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 turns out to be large, the difference is negligible. Therefore,

$$W_{set} \approx \frac{\lambda(1/\lambda^2 + \sigma_s^2/m^2)}{2(1-\rho-\frac{\lambda}{m}\beta_{TSP}\frac{\sqrt{A}}{v\sqrt{n}})}. \tag{13}$$

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

$$n > \frac{\lambda^2\beta_{TSP}^2 A}{m^2 v^2 (1-\rho)^2}. \tag{14}$$

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

The waiting time given in (13) is not the wait for service of an individual demand; it is the wait in queue for a set. The time of arrival of a set is actually the time of arrival of the last demand in that set. Therefore, to get the system time of a demand we must add to (13) 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 $\frac{n-1}{2\lambda} \leq \frac{n}{2\lambda}$. By doing the same conditioning and noting that the travel time around the tour is no more that the time to traverse the tour itself, the expected wait for service once a customer's set enters service is, for large n, no more than $\beta_{TSP}\sqrt{nA}/v + \frac{1}{n}\sum_{k=1}^{n} k\overline{s} \leq \beta_{TSP}\sqrt{nA}/v + \frac{n}{2}\overline{s}$. Therefore, if the total system time is denoted T_{GGm}, the for $\rho \to 1$,

$$T_{GGm} \leq \frac{\lambda(1/\lambda^2 + \sigma_s^2/m^2)}{2(1-\rho-\frac{\lambda}{m}\beta_{TSP}\frac{\sqrt{A}}{v\sqrt{n}})} + \frac{n(1+m\rho)}{2\lambda} + \beta_{TSP}\frac{\sqrt{nA}}{v}. \tag{15}$$

We would like to minimize (15) with respect to n to get the least upper bound. First, however, consider a change of variable to

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

Physically, y represents the ratio of the average travel time, $\overline{d} = \frac{\beta_{TSP}\sqrt{A}}{v\sqrt{n}}$ to its critical value $\frac{m(1-\rho)}{\lambda}$ (see equation (4)). With this change,

$$T_{GGm} \leq \frac{\lambda(1/\lambda^2 + \sigma_s^2/m^2)}{2(1-\rho)(1-y)} + \frac{\lambda\beta_{TSP}^2 A(1+m\rho)}{2m^2v^2(1-\rho)^2y^2} + \frac{\lambda\beta_{TSP}^2 A}{mv^2(1-\rho)y}. \tag{16}$$

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

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

Substituting this approximation into (16) and noting that for $\rho \to 1$ the approximate y^* approaches 1 we have that as $\rho \to 1$

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

The leading term is proportional to $\frac{\lambda A}{m^2v^2(1-\rho)^2}$. (By carefully following the steps above, one can verify that this term is in fact asymptotically exact.) Therefore, using Theorem 2

$$\frac{T_{GGm}}{T^*} \le \frac{m+1}{2}\frac{\beta_{TSP}^2}{\gamma^2} \qquad \text{as} \quad \rho \to 1.$$

Note that this is still dependent on m but it increase like $(m+1)/2$ rather than m as in the randomized assignment case, which is clearly better but still somewhat unsatisfactory. In the next section, we show how all dependence of the performance guarantee on m can be eliminated.

2.3.3 Independent Partitioning Policies

The last policy we examine for the uncapacitated case is to divide the region $\mathcal{A}$ into m subregions of equal size. One vehicle is assigned to each region, and it follows a single server policy μ to service demands that fall within its subregion. This will be referred to as the P_μ policy. We shall assume resulting subregions are reasonably similar and compact so that the policies that require square regions (e.g. PART and SFC) can be approximately applied.

The effect of this independent partition is to reduce *both* the area and the arrival rate by a factor of $1/m$. Thus, if we let T_{P_μ} denote the system time of the independent partition policy when policy μ is used in each subregion, it is immediate that

$$T_{P_\mu} \sim \gamma_\mu^2 \frac{(\lambda/m)(A/m)}{v^2(1-\rho)^2} = \gamma_\mu^2 \frac{\lambda/A}{m^2v^2(1-\rho)^2} \qquad \text{as} \quad \rho \to 1.$$

Comparing to the lower bound in Theorem 2 implies

$$T_{P_\mu} \le \frac{\gamma_\mu^2}{\gamma^2} \qquad \text{as} \quad \rho \to 1.$$

Thus, we see rather easily that any constant factor heavy traffic policy for the single server DTRP can readily be extended to a m server policy with the same constant factor using independent partitions.

It is tempting to infer that an optimal m-server policy can be constructed from an optimal single server policy using partitions. Unfortunately, since we do not know if there exists a single constant γ such that the lower bound in Theorem 2 is tight for all m, such a conclusion is premature; however, the idea seems highly plausible and is worth a conjecture:

Conjecture 1 *Let μ^* denote an optimal single-server DTRP policy in heavy traffic. Then P_{μ^*} is an optimal m-server policy in heavy traffic.*

3 The m-Vehicle, q-Capacity DTRP

We next examine a capacitated version of the m server DTRP. To every server we associate a depot with a fixed location in $\mathcal{A}$ with the rule that servers are allowed to use only their designated depots. Let $\mathcal{D}$ denote the set of these m depot locations. We shall allow the case where several vehicles have identical depot locations so that one could model m vehicles based out of a single location or m vehicles allocated to $k < m$ locations within this framework. The capacity constraint we consider is simply an upper bound of q on the number of customers each server can visit before being required to return to its designated depot.

Before beginning, some additional notation is needed. As before, we let i index demands according to their service order. The length of the tour containing demand i is denoted c_i and the average tour length, $\overline{c}$ is defined by $\overline{c} = \lim_{i\to\infty} E[c_i]$. Also, if the location of demand i is x_i, then the radial distance from i to the closest server, r_i, is defined as $r_i = \min_{x_o \in \mathcal{D}} \|x_i - x_0\|$ and $\overline{r} = \lim_{i\to\infty} E[r_i]$. Note also that

$$\overline{r} = E[\min_{x_o \in \mathcal{D}} \|X - x_0\|]$$

where X is a uniformly distributed location in $\mathcal{A}$.

We shall also make the assumption that each tour visits *exactly* q demands. This simplifies the analysis and seems quite reasonable for the heavy traffic case. It allows us to assert, for example, that $\overline{c} = q\overline{d}$ without worrying about questions of random incidence. This can be relaxed, but to do so introduces more technicalities than it is worth in terms of extra generality and insight.

3.1 A Heavy Traffic Lower Bound

We begin with the following lower bound:

Theorem 3 *There exists a constant γ' such that*

$$T^* \geq \gamma'^2 \frac{\lambda A(1-1/q^2)}{m^2 v^2 (1-\rho-\frac{2\lambda\overline{r}}{mvq})^2} - \frac{m(1-2\rho)}{2\lambda}.$$

<u>Proof</u>

Consider demand i and the tour of length c_i that contains it. Randomly and independently select two points in this tour with probability $1/q$ and denote them j_1 and j_2. (Note that $j_1 = j_2$ is possible as well as $i = j_1$, etc..) Also, define $j_* = \min\{j_1, j_2\}$ and $j^* = \max\{j_1, j_2\}$. Note that the length of the path from the depot to j_1 is at least r_{j_1}, since this is the distance to the closest depot. Similarly, the length of the path from the depot to j_2 is at least r_{j_2}. Adding to these two quantities the distance travel from j_* to j^* we obtain the following bound on the tour length,

$$c_i \geq r_{j_1} + r_{j_2} + \sum_{j=j_*+1}^{j^*} d_j.$$

Since the points j_1 and j_2 were randomly selected from among the q points in the tour containing i, it follows that the limiting distribution of r_{j_1} and r_{j_2} is the same as r_i. Similarly, the limiting distribution of each term d_j above is the same as d_i. Therefore, taking expectations on both sides, letting $i \to \infty$ and noting that $j_1 \to \infty$ and $j_2 \to \infty$ as $i \to \infty$ we obtain

$$\begin{aligned} \overline{c} &\geq 2\overline{r} + E[\sum_{j=j_*+1}^{j^*} d_j] \\ &= 2\overline{r} + E[|\Delta j|]\overline{d} \end{aligned} \tag{17}$$

where $\Delta j = j_1 - j_2$. The last equality above follows from the linearity of expectation and the fact that the $|\Delta j|$ is independent of the distances d_j. We next need the following lemma:

Lemma 3 $E[|\Delta j|] = \frac{1}{3}(q - \frac{1}{q})$.

<u>Proof</u>

The random variable Δj is distributed as the difference between two independent, equiprobable selections from the set $\{1, 2, \ldots, q\}$. By considering the joint sample space, it is easy to show that

$$\begin{aligned} E[|\Delta j|] &= \frac{2}{q^2}\left[1(q-1) + 2(q-2) + \cdots + q(0)\right] \\ &= \frac{2}{q^2}\left[q\sum_{i=1}^{q} i - \sum_{i=1}^{q} i^2\right]. \end{aligned}$$

Using the fact that $\sum_{i=1}^{q} i = q(q+1)/2$ and $\sum_{i=1}^{q} i^2 = q(q+1)(2q+1)/6$ and substituting above establishes the lemma.

□ (Lemma 3)

Using Lemma 3 and noting that lower bound from Lemma 1 applies in the capacitated case as well, (17) becomes

$$\bar{c} \geq 2\bar{r} + \frac{1}{3}(q - \frac{1}{q})\frac{\gamma\sqrt{A}}{\sqrt{N + m/2}}$$

Using the fact that $\bar{d} = \bar{c}/q$ and defining $\gamma' = \gamma/3$ implies

$$\bar{d} \geq \frac{2\bar{r}}{q} + \gamma'(1 - \frac{1}{q^2})\frac{\gamma\sqrt{A}}{\sqrt{N + m/2}}$$

Substituting this into the stability equation (4), rearranging and noting that $N = \lambda W$ and $T = W + \bar{s}$ we obtain the bound in Theorem 3.

□ (Theorem 3)

A few comments on this bound are in order. First note that for $q = 1$ it is trivial. This has to be true since the unit capacity case behaves like a M/G/1 queue and therefore cannot exhibit $(1-\rho)^{-2}$ behavior. Second, note that the constant value is one third of the value in the uncapacitated case. This appears to be just a by-product of the randomization used in the proof rather than a fundamental characteristic of the capacitated case. For example, if one could show that the *closest* point on the tour containing i was asymptotically an averaged distance of $\bar{r}$ from the depot, then the 1/3 factor could be eliminated. We therefore believe that the same constant holds for both the uncapacitated and capacitated cases.

Third, we see that the stability condition is now

$$\rho + \frac{2\lambda\bar{r}}{mvq} < 1.$$

As mentioned in §1, this is no longer independent of the service region geometry due to the presence of $\bar{r}$. The first term is the fraction of time required per server for on-site service while the second is the fraction of radial travel time required per server. Notice that it is only these radial travel costs that determine the stability and not the costs of the local tours. However, the growth as a function of this stability condition is precisely the same as in the uncapacitated case.

Finally, for $q >> 1$, the term $1 - 1/q^2$ is essentially one, and therefore we shall consider policies to be within a constant factor even if they lack this term.

3.2 An Optimal Light Traffic Policy

Recall that vehicles following the mSQM policy service only one customer between visits to their respective depots. This policy is therefore feasible for any capacity $q > 0$. Using the fact that the lower bound in Theorem 1 is for a relaxed problem (i.e. infinite capacity), it is therefore immediate that mSQM is also optimal for the capacitated problem in light traffic.

3.3 A Heavy Traffic Policy for the Case $m = 1$

We next construct a policy for the single server, q-capacitated problem for the case where $\rho + \frac{2\lambda\bar{r}}{vq} \to 1$ (heavy traffic). The policy will be extended to the multi-server case in the next section. It is called the qTSP policy and is defined as follows:

- **The qTSP Policy**: Divide the region $\mathcal{A}$ into n equal sized subregions (except perhaps on the boundary) using a square grid system centered at the depot as shown in Figure 1. When q consecutive demands arrive in a single subregion consider it the arrival of a set. Service sets in FCFS order as follows:

  1. Form a TSP tour on the q demands in the set.
  2. Select one of the q demands in the set at random.
  3. Service the set by traveling to the selected customer, then around the tour (servicing demands as they are encountered), and finally returning from the selected customer back to the depot.

  Optimize over the number of subregions n.

To analyze this policy we proceed as in the G/G/m case and determine the waiting time for a set, W_{set}. Let τ_i be length of the local, TSP tour containing the i-th customer and

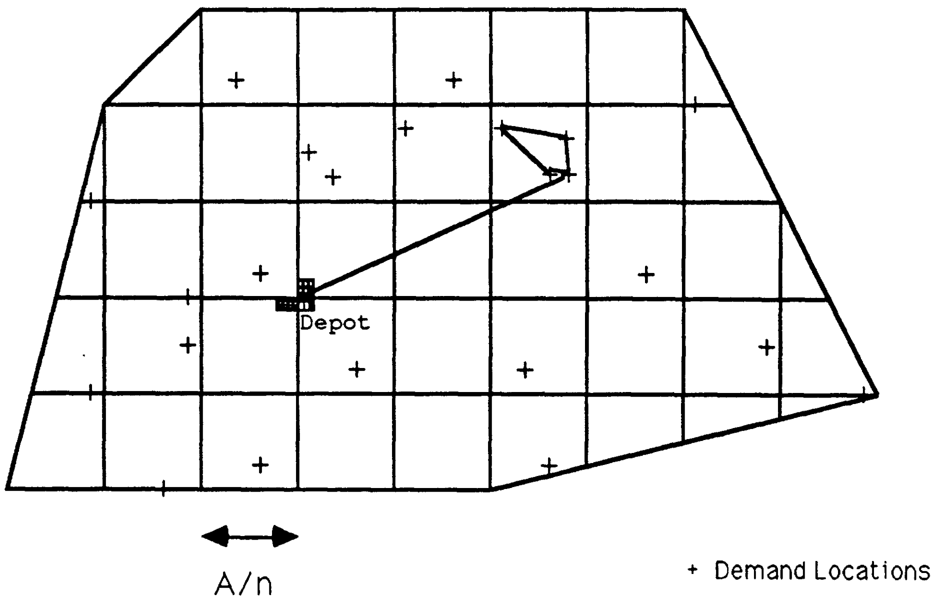

Figure 1: Subregions of the qTSP Policy

$\overline{\tau} = \lim_{i\to\infty} E[\tau_i]$. Let $\overline{r}$ be defined as above. From the uniformity of the partitions and the construction of the tours we have that the expected time to service a tour is

$$2\frac{\overline{r}}{v} + \frac{\overline{\tau}}{v} + q\overline{s}.$$

Also, denoting the ergodic variance of a random variable Z_i by σ_Z^2, the variance of the time to service a set is

$$4\frac{\sigma_r^2}{v^2} + \frac{\sigma_\tau^2}{v^2} + q\sigma_s^2.$$

We point out that σ_s^2 is assumed finite, σ_r^2 is finite due to the boundedness of $\mathcal{A}$ and σ_τ^2 is also finite (c.f. [24]).

Since the service of sets forms a G/G/1 queue, we can use these expressions in (8), which is known to be an upper bound for the case $m = 1$ [22]. This yields

$$W_{set} \leq \frac{\frac{\lambda}{q}\left(\frac{q}{\lambda^2} + \frac{\sigma_\tau^2}{v^2} + \frac{4\sigma_r^2}{v^2} + q\sigma_s^2\right)}{2\left(1 - \frac{\lambda}{q}\left(\frac{2\overline{r}}{v} + \frac{\overline{\tau}}{v} + q\overline{s}\right)\right)}$$

It is well known [20] that for q uniformly distributed points in a square region of area A that there exists a constant, $\bar{\beta}$ such that the length of the optimal tour τ satisfies

$$E[\tau] \leq \bar{\beta}\sqrt{qA}.$$

This can be shown for example using the strip heuristic, in which case a value $\bar{\beta} = 2$ is obtained. If q is large, one could reasonably use the asymptotic value $\beta \approx 0.72$. (For the reader concerned about the non-square regions on the boundary, observe that these can be considered as complete squares with a nonuniform distribution of point locations in which case the above bound still holds (see [6]).)

Since each subregion has an area A/n, substituting the bound for $\bar{r}$ into the bound on W_{set} gives

$$W_{set} \leq \frac{\lambda(\frac{1}{\lambda^2} + \frac{\sigma_r^2}{qv^2} + \frac{4\sigma_l^2}{qv^2} + \sigma_s^2)}{2(1 - \rho - \frac{2\lambda\bar{r}}{qv} - \frac{\lambda\bar{\beta}\sqrt{A}}{v\sqrt{qn}})}$$

where $\rho = \lambda\bar{s}$. By the same reasoning as in the mTSP policy, to get the system time of a demand we must add to W_{set} the expected wait for a set to form, which is at most $\frac{qn}{2\lambda}$, and the expected wait for service once the set enters service, which is at most $\frac{\bar{r}}{v} + \bar{\beta}\frac{\sqrt{qA}}{v\sqrt{n}} + \frac{q\bar{s}}{2}$. Doing this and making the change of variable

$$y = \frac{\lambda\bar{\beta}\sqrt{A}}{v(1 - \rho - \frac{2\lambda\bar{r}}{qv})\sqrt{qn}}$$

we obtain the rather complicated expression

$$T_{qTSP} \leq \frac{\lambda(\frac{1}{\lambda^2} + \frac{\sigma_r^2}{qv^2} + \frac{4\sigma_l^2}{qv^2} + \sigma_s^2)}{2(1 - \rho - \frac{2\lambda\bar{r}}{qv})(1 - y)} + \frac{\lambda\bar{\beta}^2 A}{2v^2(1 - \rho - \frac{2\lambda\bar{r}}{qv})^2 y^2} + \frac{q}{\lambda}(1 - \rho - \frac{2\lambda\bar{r}}{qv})y + \frac{q\bar{s}}{2} + \frac{\bar{r}}{v}$$

In this case, y has the interpretation as the ratio of the average *local* travel time per demand to its critical value. We can again obtain an approximate minimizing value y^* for the case $\rho + \frac{2\lambda\bar{r}}{qv} \to 1$ by linearizing the second term above about the value $y = 1$. This yields

$$y^* \approx 1 - \sqrt{\frac{v^2(\frac{1}{\lambda^2} + \frac{\sigma_r^2}{qv^2} + \frac{4\sigma_l^2}{qv^2} + \sigma_s^2)(1 - \rho - \frac{2\lambda\bar{r}}{qv})}{2\bar{\beta}^2 A}}$$

For $\rho + \frac{2\lambda\bar{r}}{qv} \to 1$, the above approximate y^* approaches 1, so asymptotically

$$T_{qTSP} \leq \frac{\lambda\bar{\beta}^2 A}{2v^2(1 - \rho - \frac{2\lambda\bar{r}}{qv})^2} + \frac{\lambda\bar{\beta}\sqrt{A}\sqrt{(\frac{1}{\lambda^2} + \frac{\sigma_r^2}{qv^2} + \frac{4\sigma_l^2}{qv^2} + \sigma_s^2)}}{v\sqrt{2}(1 - \rho - \frac{2\lambda\bar{r}}{qv})^{3/2}} + \frac{q\bar{s}}{2} + \frac{\bar{r}}{v}$$

Finally, comparing the leading term above to the bound in Theorem 3 establishes that

$$\frac{T_{qTSP}}{T^*} \leq \frac{\bar{\beta}^2}{2\gamma'^2(1 - 1/q^2)} \qquad \text{as } \rho + \frac{2\lambda\bar{r}}{qv} \to 1.$$

Hence, the policy is within a constant factor of optimum in the heavy traffic case. Note that this analysis has also established the sufficiency of the stability condition $\rho + \frac{2\lambda\bar{r}}{qv} < 1$ since T_{qTSP} is finite whenever this condition is satisfied.

The existence of such a policy also allows us to establish the following Theorem:

Theorem 4 *In the single-vehicle DTRP with vehicle capacity $q > 1$, suppose that one has the option of locating the depot anywhere within $\mathcal{A}$. Then in heavy traffic, the median is the optimal location.*

Proof

The proof is by contradiction. Suppose there exists a policy μ^* that is optimal in heavy traffic (i.e. yields the value T^* asymptotically), but it does not use the median for its depot location. Let $\bar{r}^*$ denote the expected radial distance from the median location and $\bar{r}_{\mu^*}$ denote the expected radial distance from the policy μ^* depot location. Because we have assumed policy μ^* does not use the median location, we have that $\bar{r}_{\mu^*} = \bar{r}^* + \Delta\bar{r}$ where $\Delta\bar{r} > 0$. Now consider the qTSP policy with the depot located at the median. For notational convenience, define $\delta = 1 - \rho - \frac{2\lambda(\bar{r}^* + \Delta\bar{r})}{vq}$ and $\epsilon = 2\lambda\Delta\bar{r}/vq$. By our qTSP results and Theorem 3, if μ^* is indeed optimal, then for all $\delta > 0$, T_{μ^*} must satisfy

$$\frac{\gamma'^2\lambda A(1 - 1/q^2)}{v^2\delta^2} - \frac{(1 - 2\rho)}{2\lambda} \leq T_{\mu^*} \leq \frac{\bar{\beta}^2\lambda A}{v^2(\delta + \epsilon)^2} + o(\delta + \epsilon)^{-3/2}.$$

Note, however, that for $\delta \to 0$, the lower bound above approaches infinity but the upper bound remains finite since $\epsilon > 0$. Therefore, T_{μ^*} cannot satisfy this condition for all $\delta > 0$ and hence μ^* cannot be optimal.

□ (Theorem 4)

3.4 A Heavy Traffic Policy for Some Symmetric $m > 1$ Cases

The qTSP policy can be extended to the multi-vehicle ($m > 1$) for several symmetric depot cases as shown below. These include the case where all vehicles are based out of a single depot and the case where there are $m = kp$ vehicles based out of k depots for some integer p such that the induced Voronoi cells are identical. (This includes the case where k is large and $\mathcal{D}$ is the set of k-medians)

3.4.1 The Single Depot, Capacitated m-Vehicle Case

Consider the case where the m vehicles are based out of a single depot in $\mathcal{A}$. (Or, to be consistent with our earlier convention, we say that all m depots are coincident.) We define the following G/G/m version of the qTSP policy, which we call the qG/G/m policy:

- **The qG/G/m Policy**: Divide the region $\mathcal{A}$ into a grid of n equal sized squares centered at the depot as in the qTSP policy. Form sets in exactly the same way as in the qTSP policy but serve them FCFS with the m vehicles as in a G/G/m queue. Optimize over n.

The analysis of this policy proceeds exactly as for the qTSP except that the general version of the G/G/m limit (8) is used. We will not repeat the analysis since it tedious and follows almost exactly that of the qTSP policy. The result is that

$$T_{qGGm} \leq \frac{\lambda(\frac{1}{\lambda^2} + \frac{\sigma_r^2}{qv^2} + \frac{4\sigma_r^2}{qv^2} + \sigma_s^2/m^2)}{2(1-\rho-\frac{2\lambda\bar{r}}{qvm})(1-y)} + \frac{\lambda\bar{\beta}^2 A}{2m^2v^2(1-\rho-\frac{2\lambda\bar{r}}{qvm})^2y^2} + \frac{qm}{\lambda}(1-\rho-\frac{2\lambda\bar{r}}{qvm})y + \frac{q\bar{s}}{2} + \frac{\bar{r}}{v}$$

where y is defined in terms of n by

$$y = \frac{\lambda\bar{\beta}\sqrt{A}}{vm(1-\rho-\frac{2\lambda\bar{r}}{qvm})\sqrt{nq}}$$

An approximate optimal y is given by

$$y^* \approx 1 - \sqrt{\frac{v^2(\frac{1}{\lambda^2} + \frac{\sigma_r^2}{qv^2} + \frac{4\sigma_r^2}{qv^2} + \sigma_s^2/m^2)(1-\rho-\frac{2\lambda\bar{r}}{qvm})}{2\bar{\beta}^2 A}}.$$

For $\rho + \frac{2\lambda\bar{r}}{qvm} \to 1$, this value of y approaches 1 and therefore by substitution into the expression for T_{qGGM} we obtain

$$T_{qGGm} \sim \frac{\lambda\bar{\beta}^2 A}{2m^2v^2(1-\rho-\frac{2\lambda\bar{r}}{qvm})^2}, \qquad \text{as } \rho + \frac{2\lambda\bar{r}}{qvm} \to 1,$$

where the above asymptotics are interpreted to be upper bounds if the constant $\bar{\beta}$ is only a bound. Comparing to the bound in Theorem 3, we see that

$$\frac{T_{qGGm}}{T^*} \leq \frac{\bar{\beta}^2}{2\gamma'(1-1/q^2)} \qquad \text{as } \rho + \frac{2\lambda\bar{r}}{qvm} \to 1,$$

and thus the qG/G/m policy has a constant factor performance guarantee in heavy traffic that does not depend on m.

An exact analogue of the argument used in proving Theorem 4 gives us the following theorem for the single-depot case (we omit the proof):

Theorem 5 *Consider the m-vehicle DTRP where each vehicle has capacity $q > 1$ and all vehicles operate out of the same depot location. Suppose that one has the option of locating this depot anywhere within $\mathcal{A}$. Then in heavy traffic, the median is the optimal location.*

Just as in the uncapacitated case, randomly assigning demands to vehicles (randomization) produces a bound proportional to m. This can be seen by observing that randomization produces a system time equal to T_{qTSP} with the arrival rate reduced by a factor of $1/m$. This may seem odd at first since both qG/G/m and the randomization policy construct tours in exactly the same fashion; however, with randomization, m different *types* of sets are formed within each square, one type for each of the m servers. Thus, sets take m times as long to form as in the qG/G/m policy, and this accounts for the extra factor of m in the system time.

Some readers may question why one cannot simply let $q \to \infty$ in the above bound and obtain an improvement over the earlier uncapacitated result. The problem with this is that the term relating to the wait for service once a demand's set enters service is proportional to q. Indeed, it is precisely this term that accounts for the additional $m/2$ factor for the uncapacitated case.

Finally, we point out that an independent partition policy for this case is possible if the region $\mathcal{A}$ is a circle with the depot at its center. If this is true, simply dividing the circle into m slices of equal area and applying a single server, qTSP policy in each slice produces the same constant factor guarantee as the G/G/m policy. This is obtained by using the qTSP results and observing that the conditional distribution of r_i is each slice is the same as the unconditional distribution, the area is A/m, the arrival rate is λ/m and the expected system time in each slice is identical. This policy, however, should be considered inferior since it only applies to the symmetric case whereas the qG/G/m policy applies to any region $\mathcal{A}$ with an arbitrary depot location.

3.4.2 The Case of k Symmetric Depot Locations

We now briefly describe a multi-depot case for which provably good policies can be constructed. Suppose there are k depots and a positive integer p such that $m = kp$. That is, there are exactly p vehicles per depot. Further, suppose these k depots induce Voronoi cells that are identical in shape and size. Then if one applies a p vehicle policy (i.e. qG/G/p) in each cell, the resulting system time will be within a constant factor of the lower bound in heavy traffic. This due to the fact that each cell has an arrival rate of λ/k and serves an area of size A/k, each of which has the same mean radial distance $\bar{r}$. Therefore, since each region operates with p vehicles we have

$$\begin{aligned} T &\sim \frac{\bar{\beta}^2(\lambda/k)(A/k)}{p^2v^2(1-\rho-\frac{2(\lambda/k)\bar{r}}{vqp})^2}, \qquad \text{as } \rho+\frac{2\lambda\bar{r}}{qvm}\to 1 \\ &= \frac{\lambda\bar{\beta}^2 A}{m^2v^2(1-\rho-\frac{2\lambda\bar{r}}{vqm})^2}, \end{aligned}$$

and hence the policy has a constant factor performance guarantee.

If k is large and the depots are located at the k median locations, then Haimovich and Magnanti [17] show that the Voronoi cells approach a uniform, hexagonal partition of A (i.e a honeycomb pattern). Since this simultaneously produces uniform Voronoi cells and minimizes $\bar{r}$, one can show that assigning p vehicles to each of the k medians is again provably good.

In the asymmetric case, it is less clear what approach to take. Certainly if $m = kp$ and one has the option of positioning depots, then some approximately uniform partition seems best. If the depot locations are fixed at asymmetric locations and/or the m vehicles cannot be evenly partitioned among the depot locations, then it is less certain which policy is best. Indeed, there seems to be an inherent contradiction in the asymmetric case: each set must be serviced by its closest depot to achieve a radial travel cost of $\bar{r}$ yet the arrivals must be evenly allocated to vehicles to achieve a uniform rate of λ/m. As a result, for these cases there appears to be no way to construct provably good policies relative to the lower bound in Theorem 3. More sophisticated bounds are no doubt needed.

4 Extensions to Higher Dimension and General Spatial Distribution

The above results can be extended to Euclidean subsets $\mathcal{A}$ of $\Re^d$ and to the case where demand locations have a nonuniform distribution. These extensions are explained briefly below.

4.1 Bounds and Policies for Higher Dimensions

Let V denote the volume of $\mathcal{A} \in \Re^d$. Then repeating the proof of Theorem 2 for general d one can show the following theorem for the uncapacitated, m-server DTRP.

Theorem 6

$$T^* \geq \gamma(d)^d \frac{\lambda^{d-1} V}{v^d m^d (1-\rho)^d} - \frac{m(1-\rho)}{2\lambda}$$

where

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

and $c_d = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2}+1)}$ *is the volume of a ball of unit radius in* $\Re^d$.

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

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

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

For example, the P_{TSP} policy (i.e. the partition policy using a single-server TSP in each of the m subregions) in d dimensions has a constant of $\gamma_{mTSP}(d) = \beta_{TSP}(d)$, where $\beta_{TSP}(d)$ is the d dimensional TSP constant.

An interesting result is found by examining this policy for $d \to \infty$. In [11], it is conjectured that for $d \to \infty$

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

This is based on similar results for the minimum spanning tree and minimum matching problem. By using the fact that for $d \to \infty$, $\frac{d}{d+1} \sim 1$, $\left(\frac{1}{2(d+1)}\right)^{1/d} \sim 1$ and $\Gamma(\frac{d}{2}+1) \sim \sqrt{2\pi}(\frac{d}{2})^{\frac{d}{2}+\frac{1}{2}} e^{-\frac{d}{2}}$, it is straightforward to show that

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

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

Theorem 7 *Consider the uncapacitated, m server DTRP. If* $\beta_{TSP}(d) \sim \frac{\sqrt{d}}{\sqrt{2\pi e}}$, *then* P_{TSP} *is an optimal heavy traffic policy asymptotically as* $d \to \infty$.

4.2 General Spacial Distribution

Suppose the locations of demands are not uniformly distributed but are distributed according to a general distribution with bounded support and absolutely continuous part $f(x)$. Then it is well known [32] that the length of a TSP tour on n points drawn independently from this distribution is asymptotic ($n \to \infty$) to

$$\beta_{TSP}(d) n^{\frac{d-1}{d}} \int f(x)^{\frac{d-1}{d}} dx$$

where $\beta_{TSP}(d)$ is the d-dimensional TSP constant mentioned above.

Using this fact, one can show that the the P_{TSP} policy for the m-server, uncapacitated DTRP has a system time that satisfies

$$T_{P_{TSP}} \sim \gamma_\mu(d)^d \frac{\lambda^{d-1}\left(\int f(x)^{\frac{d-1}{d}} dx\right)^d}{v^d m^d (1-\rho)^d} \qquad \text{as } \rho \to 1.$$

For $f(x) = \frac{1}{V}$, this reduces to our previous result. Unfortunately, due to the non-uniformity, the techniques used in the lower bound of Theorem 2 cannot be directly applied in this case; however, we conjecture that an analogous result holds; namely,

Conjecture 2 *For the uncapacitated, m-server DTRP, there exists a constant $\gamma(d)$ such that*

$$T^* \geq \gamma(d)^d \frac{\lambda^{d-1}\left(\int f(x)^{\frac{d-1}{d}} dx\right)^d}{v^d m^d (1-\rho)^d} - \frac{m(1-\rho)}{2\lambda}$$

For the capacitated case, we conjecture that a lower bound of this form above holds with $(1-\rho)$ replaced by $(1 - \rho - \frac{2\lambda\bar{r}}{mvq})$ and $\gamma(d)$ replaced by $\gamma(d)/3$.

5 Optimization Problems

In most vehicle routing systems, there are three major resource allocation and/or operational decisions that must be made:

- **Select the Fleet Composition** - This decision involves choosing both the number and type of vehicles to deploy within the service region.
- **Assign Vehicles to Districts** - Once the composition of the fleet has been decided, the individual vehicles must be allocated to various subregions or subsets of customers (*districts*).

- **Route Vehicles within Districts** - Given an assignment of vehicles to districts, routes must be found for the individual vehicles that minimize travel cost, waiting time or some combination of these costs.

These decisions form a natural and integrated hierarchy. The highest level is fleet composition which is a long-term strategic decision that requires estimates of how effectively vehicles can be apportioned and routed. Districting is an intermediate-term decision that is made based on knowledge of the fleet composition as well as estimates of the cost of routing within a given district. Finally, the short-term, tactical routing decisions require knowledge of both the type of vehicle and the district or customers to be served, which are provided by the two higher levels.

In this section, we give some brief examples of how DTRP models can be used within this hierarchy of allocation/operations decisions. This collection of problems is not meant to be exhaustive. Rather, it is intended to suggest how our results could be used as building blocks for some insightful normative models. The reader will no doubt think of many additional problems in this vein.

5.1 Optimal Fleet Composition/Sizing

Consider the following strategic, fleet composition problem: a utility firm would like to acquire a fleet of m repair vehicles each of capacity q to service its network. The vehicles are to be based out of a single depot. The objective is tc minimize *total* operating cost, which is a linear combination of the downtime (system time) cost, $c_1 T(m)$, and the vehicle operating costs (depreciation, wages, fuel, etc.), $c_2 m$; that is,

$$\min_m c_1 T(m) + c_2 m.$$

Suppose failures are quite frequent so it is decided that the qG/G/m policy is to be used. In this case, an approximate expression for the average downtime is

$$T(m) \approx \frac{\lambda \bar{\beta}^2 A}{2m^2 v^2 (1 - \rho - \frac{2\lambda \bar{r}}{qvm})^2}.$$

Substituting this into the minimization above and ignoring integrality, it is easy to find that the optimal m is

$$m^* = \left(\frac{2c_1 \lambda \bar{\beta}^2 A}{c_2 v^2} \right)^{1/3} + \lambda \bar{s} + \frac{2\lambda \bar{r}}{vq}.$$

(One would of course round up or round down this solution to achieve the best integer m.)

Observe that the first term is simply the amount by which m exceeds its critical value $\lambda\bar{s} + \frac{2\lambda\bar{r}}{vq}$ and that this excess increases and decreases with the various system parameters and costs as expected. Also note that with the lower level decisions are implicit in the formulation; namely, route vehicles using the qG/G/m policy and assign customers evenly to each of the m vehicles.

5.2 Optimal Districting

An example in this category is the following somewhat stylized districting problem: consider a company that has a fleet of m *heterogeneous* vehicles with velocities $v_1, \ldots, v_m$ each of unlimited capacity. For instance, the fleet might consist of k older slow vehicles and $m-k$ newer fast vehicles. We would like to assign each vehicle a portion of the service region so as to minimize the average system time. We shall assume the system operates in heavy traffic and that each vehicle follows a heavy traffic policy μ in its assigned region.

If the fraction of service area assigned to vehicle i is denoted p_i, then the optimization problem is

$$\begin{aligned} &\min_{s.t.} \sum_{i=1}^{m} \frac{\lambda\gamma_\mu^2 A p_i^3}{v_i^2(1-\lambda\bar{s}p_i)^2} \\ &\sum_{i=1}^{m} p_i = 1 \\ &p_i \geq 0 \quad i = 1, \ldots, m. \end{aligned}$$

If we introduce a Lagrange multiplier, μ, on the equality constraint and define the functions

$$f_i(p) = \frac{\lambda\gamma_\mu^2 A p^3}{v_i^2(1-\lambda\bar{s}p)^2},$$

then the optimal p_i^* and μ^* satisfy

$$\begin{aligned} \dot{f}_i(p_i^*) &= \mu^* \\ \sum_{i=1}^{m} p_i^* &= 1 \\ p_i^* &\geq 0 \quad i = 1, \ldots, m. \end{aligned}$$

This is complicated to solve analytically in the general case, but it is not difficult numerically.

A simplification occurs for the case $\bar{s} = 0$ (and hence $\rho = 0$), $\sigma_s^2 = 0$ and $\lambda \to \infty$. That the above expressions f_i remain valid in this case has not been demonstrated; however, by reexamining the arguments for the G/G/m policy for this case, it is easy to verify that the system time is indeed given by the expressions f_i above. For this case, $f_i(p) = \frac{\lambda\gamma_\mu^2 A p^3}{v_i^2}$ and the optimal solution is

$$\begin{aligned} p_i^* &= \frac{v_i}{\sum_{i=1}^m v_i} \\ \mu^* &= \frac{3\gamma_\mu^2 \lambda A}{(\sum_{i=1}^m v_i)^2}, \end{aligned}$$

which can be verified by substitution into the optimality conditions above. Thus, it is optimal to allocate the area proportional to vehicle velocities in the special case where on-site service times are negligible.

Note that in this example, we assume the fleet composition decision has already been made and that the routing within each district will use the TSP policy.

5.3 Routing to Minimize Travel and Waiting Cost

Choosing routes for individual vehicles is the lowest level decision in our hierarchy, and making these decisions so as to minimize system times has been the primary focus of this paper. Yet in many practical problems, there is a mixed objective of minimizing waiting ***and*** travel costs. In our discussion thus far, travel cost has largely been ignored; however, it turns out that a travel cost is indeed present, and it plays an important role in the optimization of parametric policies like G/G/m and qG/G/m.

Recall that in these policies we made a change of variable from the set size n to a variable y that represented the ratio of travel time per demand to some critical value. In the uncapacitated case, y is simply the ratio of $\overline{d}/v$ to its critical value $\frac{m(1-\rho)}{\lambda}$; in the capacitated case, it is the ratio of the *local* travel cost to its critical value. Rather than seeking the y that minimizes the system time, it is useful to examine the system time as a function of y; that is, $T(y)$. Note that for $y = 0$ no traveling occurs while $y = 1$ implies the maximum amount of travel per arrival. For simplicity, we shall concentrate on the case of a single, uncapacitated vehicle (i.e. the TSP policy as defined in the §1.3) to illustrate the tradeoff. Similar results apply for the other cases.

In the uncapacitated, $m = 1$ case, we obtained a system time of the form

$$T(y) = \frac{c_1}{(1-y)^2} + \frac{c_2}{y^2} + \frac{c_3}{y} + c_4$$

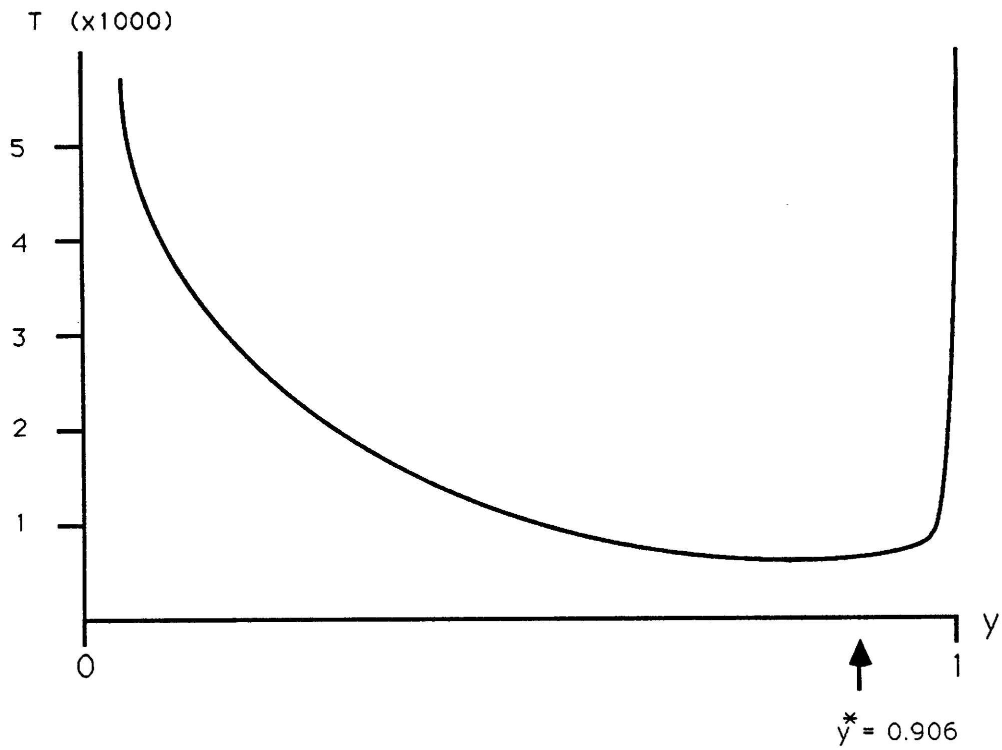

Figure 2: System Time vs. Travel Cost per Demand

where $c_1, \ldots, c_4$ depended on the system parameters. This function is shown graphically in Figure 2 for the case $\sigma_s^2 = 0$, $A = 1$ and $\overline{s} = 0.1$ and $\rho = 0.9$. Note that the function has poles at both 1 (travel equal to its critical value) and 0 (no travel at all) as expected.

To minimize $T(y)$, we want to optimally balance between these two extremes. For $\rho \longrightarrow 1$, the coefficient c_2 increases much more rapidly that c_1 and c_3. Thus, the optimal value of y approaches 1 corresponding to the travel time per customer approaching its maximum (critical) value. Note that increasing y beyond y^* increases *both* the travel cost *and* the waiting time; therefore, there is no reason to choose a value in this range. However, one might want to choose a lower value of y, corresponding to less travel per demand, at the cost of increasing the average system time. For instance, in our example $y^* \approx 0.906$ and the system time is 578. If we decide to reduce the average travel cost per demand 10% to $y = 0.81$, the system time increases by 21% to 702. One could of course place costs on *both* the travel time per demand and the system time and select a set size that minimizes their sum. The result would be to form larger sets which improve travel efficiency at the expense of increased system time.

Similar relationships are found for the capacitated case (qG/G/m policy) except that the variable y represents the ratio of the local travel cost to its critical value. The interesting difference here is that the radial costs per service, $2\lambda\overline{r}/qm$, cannot be traded-off against system time; only the local costs can. Finally, we point out that in the space filling curve (SFC) and nearest neighbor (NN) policies, which are nonparametric, no such tradeoff can be achieved. Simulation results show that they operate in heavy traffic so that $\overline{d}$ is close to its critical value. Although these policies have advantages for systems where exact parameter estimation is difficult or parameters change frequently, this lack of travel/waiting cost control should be considered a disadvantage.

6 Concluding Remarks

We have examined congestion effects when operating many capacitated vehicles in a dynamic and stochastic environment. In the uncapacitated case we found that the stability condition is independent of any characteristics of the service region while in the case where each vehicle has capacity $q < \infty$, the depot location and system geometry strongly influences the stability condition. We also showed that the distributed character of the system gives rise to behavior very different than that of traditional queues. In particular, the optimal, expected system time in heavy traffic is $\Theta(\frac{\lambda A}{m^2v^2(1-\rho)^2})$ for the uncapacitated case and $\Theta(\frac{\lambda A}{m^2v^2(1-\rho-\frac{2\lambda\overline{r}}{mvq})^2})$ for the capacitated case. Moreover, we found optimal policies in light traffic and several policies that have system times within a constant factor of the optimal policy in heavy traffic. We then extended our analysis to higher dimensions and to arbitrary demand distributions.

We also showed how our models can be used to make various resource allocation decisions. For example, although our analysis focused on system times, it turns out that there are interesting tradeoffs between travel cost and waiting time costs in several of our proposed policies.

These results give new insights into the problems of stability, fleet sizing, districting, depot location and response time under congestion for dynamic, stochastic vehicle routing systems. However, some open questions still remain in this area. In particular a proof of Conjecture 1 is needed to round out our understanding of the relationship between the single and multiple vehicle problem. Also, a proof of Conjecture 2 would formalize the generalization to arbitrary distributions, which is an important extension in practical settings. A more challenging problem is to try and close the gap between the lower bound constant γ and the various policy constants γ_μ, with the ultimate goal of finding asymptotically optimal policies in heavy traffic. However, this probably requires a new approach to the proof of Lemma 1, which is already quite intricate. A challenging problem in a different direction is to investigate dynamic routing in a network environment rather than under some Euclidean metric. We hope that some of the insights and techniques presented in this paper can be used for this problem.

References

[1] Afrati, F. S., Cosmadakis, S. (1986), Papadimitriou, C., Papgeorgiou, G. and Papakostantinou, N., "The Complexity of the Traveling Repairman Problem", *Information Theory Applications* (France), 20, 79-87.

[2] Agnihothri, S. R. (1988), "A Mean Value Analysis of the Traveling Repairman Problem", *IIE Transactions*, 20, 2, 223-229.

[3] Baker, J. E. and Izhak, R. (1987), "Polling with a General-Service Order Table", *IEEE Transactions of Communications*, COM-35, 3, 283-288.

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

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

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

[7] 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.

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

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

[10] Bertsimas, D., Jaillet, P. and Odoni, A. (1988) "A Priori Optimization", December 1988, to appear in *Operations Research.*

[11] Bertsimas, D. and van Ryzin, G. (1989), "An Asymptotic Determination of the Minimum Spanning Tree and Minimum Matching Constants in Geometrical Probability", to appear in *Operations Research Letters.*

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

[13] Brown, G. G. and Graves, G. W. (1981), "Real-Time Dispatching of Petroleum Tank Trucks", *Management Science*, 27, 1, 19-32.

[14] Browne, S. and Yechiali, U. (1989), "Dynamic Priority Rules for Cyclic-Type Queues", *Advances in Applied Probability*, 21, 432-450.

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

[16] Golden G. and Assad A., (1988), *Vehicle Routing; Methods and Studies*, North Holland, Amsterdam.

[17] 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.

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

[19] 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.

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

[21] Kingman, J.F.C. (1964), "The Heavy Traffic Approximation in the Theory of Queues", in W. L. Smith and R. I. Wilkinson, eds., *Proceedings of the Symposium on Congestion Theory*, Univ. of North Carolina, Chapel Hill, 137-169.

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

[23] Köllerström, J. (1974), "Heavy Traffic Theory for Queues with Several Servers", *Journal of Applied Probability*, 11, 544-552.

[24] 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.

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

[26] 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*)

[27] Powell, W. (1985), "An Operational Planning Model for the Dynamic Vehicle Allocation Problem with Uncertain Demands", Princeton University, Department of Civil Engineering Working Paper EES-85-15.

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

[29] Psaraftis, H., Orlin, J., Bienstock, B. and Thompson, P. (1985), "Analysis and Solution Algoritms of the Sealift Routing and Scheduling Problems: Final Report", Working Paper No. 1700-85, Sloan School of Management, MIT.

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

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

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

[33] Stecke, K. E. and Aronson, J. E. (1985), "Review of Operator/Machine Interference Models", *International Journal of Production Research*, 23, 129-151.

[34] Takagi, H. (1986), *Analysis of Polling Systems*, The MIT Press, Cambridge, Mass..