Stochastic and Dynamic Vehicle Routing in Euclidean Service Regions

by

Garrett John van Ryzin

B.S.E.E. Columbia University
(1985)
S.M. Massachusetts Institute of Technology
(1987)

Submitted to the Operations Research Center
in Partial Fulfillment of
the Requirements for the Degree of
DOCTOR OF PHILOSOPHY IN OPERATIONS RESEARCH

at the

Massachusetts Institute of Technology

June 1991

Signature of Author ______________________________
Operations Research Center
May 3, 1991

Certified by ______________________________
Dimitris J. Bertsimas
Assistant Professor of Management Science
Thesis Supervisor

Accepted by ______________________________
Amedeo R. Odoni
Codirector, Operations Research Center

Stochastic and Dynamic Vehicle Routing in Euclidean Service Regions

by

Garrett John van Ryzin

Submitted to the Operations Research Center
on May 3, 1991
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY IN OPERATIONS RESEARCH

Abstract

We introduce and analyze a model for stochastic and dynamic vehicle routing in which vehicles traveling at a 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 begin by examining the case where the regions is served by a single, uncapacitated vehicle and demand locations are uniformly distributed. We find necessary and sufficient conditions for the existence of a stable policy, a policy that is optimal in light traffic and several policies that have system times within a constant factor of the optimum in heavy traffic. We then extend our analysis the problem in several directions. First, we analyze 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. We then consider the case in which each vehicle can serve at most q customers before returning to a depot. In contrast to the uncapacitated case, we show the stability condition in this case depends strongly on the geometry of the region. Finally, we examine the problem for generally distributed demand locations. Policies that have system times within a constant factor of the optimum in heavy traffic are proposed and analyzed for all these various extensions. Several other extensions to more general cost structures, strategic planning problems and generalized removal rules are also discussed.

Thesis Supervisor: Dimitris J. Bertsimas

Title: Assistant Professor of Management Science

Acknowledgement

I gratefully acknowledge the support provided by the National Science Foundation under grant DDM-9014751 and a fellowship provided by the MIT Center for Transportation Studies with support from the UPS Foundation. The Operations Research Center also provided much needed coordination and arranged various teaching assistantships during the course of my studies, for which I am grateful. I would especially like to thank Marcia and Paulette for helping me overcome the numerous (and often self inflicted) bureaucratic obstacles that cropped up throughout my program. I also sincerely appreciate the friendship and comradery of the students at the OR center, both past and present, and hope that we continue are friendships far into the future. (Even past the age when we can no longer play soccer!)

I would like to thank my committee members, Amedeo Odoni and Larry Wein, for their enthusiasm, support and friendship. Not only did they provide a range of useful comments and objective feedback, but they were also extremely nice people to work with! Thanks guys!

Special thanks are due to my advisor and good friend, Dimitris Bertsimas. He was always enthusiastic about our various research problems and any ideas I had about them. Indeed, I am quite certain I would not have attempted such an ambitious project without his encouragement, and frequently throughout the course of our research his confidence in my abilities exceeded mine. I also learned from him, both directly and by observation, how to be more precise and scientific in my thinking and approach to problems, which to me has been an invaluable lesson. But most importantly, I appreciate the friendship he has shown me over these past three years, for this made our collaboration truly enjoyable.

To Mary Beth and Stephanie with love and thanks.

Contents

List of Figures

Chapter 1

Introduction

1.1 Motivation

Vehicle routing problems (VRPs), in particular the classical traveling salesman problem (TSP), are some of the most studied problems in the operations research and applied mathematics literature. The attention they receive is due in large part to their richness and inherent elegance. The TSP in particular is, in many ways, the prototypical hard combinatorial problem. However, VRPs are also encountered frequently in practical distribution systems, both directly and as subproblems, and these practical applications have been a further stimulus for research.

Classically, VRPs are viewed as static, deterministic problems. A set of known customer locations defines an instance, and the objective is to visit customers so as to minimize the total travel cost, perhaps subject to certain constraints (*e.g.* a limit on vehicle capacity). This classical paradigm has generated significant research interest over the years (see for example [34], [18]) resulting in major contributions in the areas of combinatorial optimization, the analysis of heuristics and complexity theory. Yet, in many of the practical applications in which VRP's arise (*e.g.* emergency service, inventory resupply, mobile repair and distribution), their is a significant dynamic component to the problem. Thus, the classically defined VRP's are often deterministic, static approximation to practical problems which are, in reality, often both probabilistic and time varying (dynamic).

For example, a typical application of a TSP is in routing a delivery vehicle from a central depot to a set of dispersed demand points so as to minimize the total travel (delivery) costs. In real distribution systems, 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 level) if often as important as the travel cost. And in some cases, in particular emergency service, waiting time is the most important measure of system performance.

As a canonical example of a logistics application with strong probabilistic and dynamic components, consider the following utility repair problem: A utility firm (electric, gas, water and sewer, highway, etc.) is responsible for maintaining 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. Routing decisions are made based on a realtime log of current failures and perhaps some characterization of the future failure process. 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 other closely related problems to this canonical example that arise in practice. For instance, consider a firm that delivers a product from a central depot to customers based on orders that arrive in realtime. Orders are entered into a log and delivery vehicles are dispatched with the objective of minimizing some combination of the delivery cost and the average wait for delivery. Such an order process is likely to be found in firms that serve a large population of customers (or potential customers) each of whom orders relatively infrequently (e.g. home heating oil distributors, mail order firms, etc.).

Further important examples are found in finished goods distribution and freight consolidation. Consider, for example, an automobile manufacturer. Cars are produced at an assembly plant and put into finished goods inventories (parking lots) to await distribution by a fleet of car-hauling trucks. Each car is designated for a particular dealer. Conceptually, the inventory can be thought of as a "log" of locations that must be visited by the delivery vehicles. New entries to this log are made every time a new vehicle is added to the inventory, and entries are deleted when automobiles are delivered to their designated dealers. For a fixed production rate, minimizing the waiting time in this case is, by Little's theorem, equivalent to minimizing the inventory of finished goods either on the lot or in transit.

Similar distribution problems are found in freight consolidation (e.g. less-than-truckload (LTL) shipping) and parcel post systems. Here distribution centers receive partial loads designated for specific locations in a service region. These partial loads are queued (stored in a distribution terminal) and eventually consolidated into full truckloads for delivery. Lowering the wait for delivery in these systems is important both for improving the service level (delivery time) and for reducing inventory costs (terminal space, insurance costs, etc.).

It is an unfortunate fact that current models and techniques have little to say about vehicle routing when stochastic and dynamic elements are included [39] . This is due in large part to the inherent difficulties of combining vehicle routing and congestion models. In particular, including a time element along with randomness usually destroys the combinatorial structure required for classical vehicle routing methods. Similarly, the strong dependencies present in travel times usually violate the assumptions required to apply traditional queuing models. Indeed, Psaraftis [39] points out that although congestion (queueing) and vehicle routing theory are both very rich subjects, little work has been done to combine them.

Our goal in this thesis is to begin to fill this void. That is, to use techniques from established areas and to develop new techniques in order to form a useful theory of vehicle routing under congestion. We believe the results obtained represent a step toward this goal. They significantly extend the range of logistics problems that can be analyzed quantitatively. In addition, they give strong insights into the behavior of stochastic and dynamic vehicle routing systems, as well as providing a variety of practical, provably good heuristics. We hope these insights will ultimately lead to significant improvements in the design and operation of real-world logistics systems.

1.2 Literature Review

Stochastic and dynamic elements of the type mentioned above are sometimes incorporated in the classical vehicle routing framework through the use of rolling horizon procedures. These procedures 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 [38] and Psaraftis [40] for examples of this approach. Though useful for data-intensive tactical problems, they are inherently *ad hoc* and do not give insight into the fundamental behavior of these systems.

To understand such behavior more directly, 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 $(1, 2, \dots, n, 1)$ 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 [42] and Afrati *et al.* [1]). The TRP even seems difficult on trees. Minieka [35] 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 [23], Bertsimas [9], [10] and Bertsimas, Jaillet and Odoni [11] 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 realtime 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 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 examples 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 [47] 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 [16]. 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 is 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 [43] 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 applications is the dynamic traveling salesman problem (DTSP) proposed by Psaraftis [39]. 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.

1.3 Definition

The problem we investigate, which we call the dynamic traveling repairman problem (DTRP) is a Euclidean model of a dynamic VRP. Demands for service arrive according to a Poisson process with rate λ to a connected, bounded Euclidean service region $\mathcal{A}$ of area A. Upon arrival, demands assume an independent and identically distributed (i.i.d.) location in $\mathcal{A}$ according to a continuous density $f(x)$ defined over $\mathcal{A}$. Demands are serviced by m identical vehicles that travel at constant velocity v. At each location, vehicles spend some time s in on-site service that is i.i.d. and generally distributed with finite first and second moments denoted by $\overline{s}$ and $\overline{s^2}$ respectively. Figure 1.1 shows a picture of this system.

Initially, we shall assume that there are no capacity constraints on the vehicles. In Chapter 4 we consider the case where there is an upper bound of q on the number of demands that can be served before a vehicle must return to a designated depot location.

A policy for routing the vehicles is called stable if the number of demands in the system is bounded almost surely for all times t. Let $\mathcal{M}$ denote the set of stable policies. If a policy is stable, $\rho \equiv \frac{\lambda \overline{s}}{m}$ is the fraction of time vehicles spend in on-site service. We write T_μ to indicate the system time of a particular policy $\mu \in \mathcal{M}$. The DTRP is then defined as the optimization problem

DTRP:

$$\min_{\mu \in \mathcal{M}} T_\mu .$$

We let T^* denote the optimal value in this minimization.

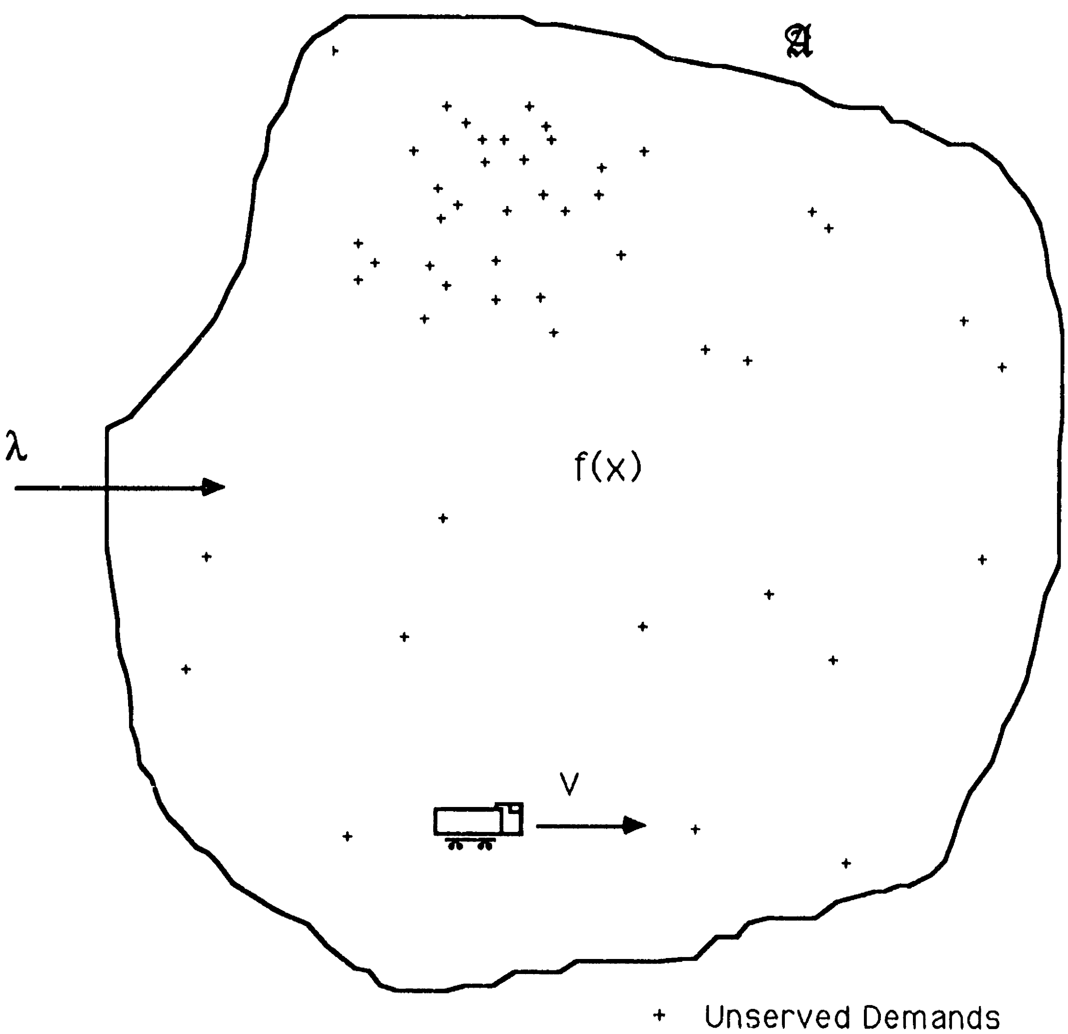

Figure 1.1: Pictorial Representation of the DTRP

1.4 Overview and Outline of Thesis

1.4.1 The Single Uncapacitated Vehicle DTRP

Our analysis of this problem requires several known results from from queuing theory and geometrical probability. Chapter 2 lists these results and gives appropriate references.

In Chapter 3 we analyze the DTRP for the case where the entire region is served by a single uncapacitated vehicle. In the light traffic case ($\lambda \to 0$) we show that a policy based on locating the server at the median, x^*, of $\mathcal{A}$ and serving customers in FCFS order, returning to the median after each service is optimal. The optimal expected system time, T^*, in this case satisfies

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

where $X \sim f(x)$. (Note that the first term above is simply the expected travel time from the median). We extend this result to the m-vehicle case in Chapter 4, in which case the first term above becomes the m-median distance divided by the velocity.

We then investigate the problem in heavy traffic, adding the restriction that demands are uniformly distributed ($f(x) = 1/A$). We show that for this single, uncapacitated vehicle problem policies exist that have finite system times for all $\rho < 1$. This is surprising in that the condition is independent of the service region size and shape; it is also the mildest stability restriction one could hope for. We then show that there exists a constant $\gamma = \frac{2}{3\sqrt{2\pi}} \approx 0.266$ such that

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

We extend this bound to the m-vehicle case in Chapter 4. An analogous asymptotic bound is derived in Chapter 5 that improves the constant by a factor of $\sqrt{2}$ to $\gamma = \frac{2}{3\sqrt{\pi}} \approx 0.376$. Note that this bound grows like $(1-\rho)^{-2}$ as $\rho \to 1$. Thus, though the stability condition is similar to that of a traditional queue, the system time increases much more rapidly as congestion increases.

In Chapter 3 we propose several policies μ that have finite system times, T_μ, for all $\rho < 1$. In addition, we show that these policies have the same asymptotic behavior, namely

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

where the constant γ_μ demands only on the policy μ. Hence, by comparing this to the lower bound above, we see that the ratio T_μ/T^* is bounded as $\rho \rightarrow 1$. (Such a bound is henceforth called a *constant factor guarantee.*) The provably best policy is one based on forming optimal traveling salesman tours for which $\gamma_\mu = \beta/\sqrt{2} \approx 0.51$, where β is the Euclidean TSP constant [6]. Relative to the value $\gamma = \frac{2}{3\sqrt{\pi}} \approx 0.376$ from Chapter 5, this gives our best provable guarantee of

$$\frac{T_\mu}{T^*} = \frac{\beta^2}{2\gamma^2} \approx 1.8.$$

We conjecture that in fact the TSP policy is optimal in heavy traffic and thus the ratio above is one.

Using simulation, we also analyze a policy based on the space filling curve heuristic [4] and one based on serving the nearest neighbor. The results show a similar $\gamma_\mu^2 \frac{\lambda A}{v^2(1-\rho)^2}$ behavior, even for moderate value of ρ. We estimate the constants γ_μ for these policies from the simulation results. Though not competitive with the best TSP policy, these heuristics are simple to implement and of very low complexity, which may make them attractive in practice.

1.4.2 The Multiple Capacitated-Vehicle DTRP

As satisfying as these results are, the model of a single uncapacitated vehicle is somewhat unrealistic for most practical purposes. Therefore, in Chapter 4 we extend our results to the case where demands are still uniformly distributed, but the region $\mathcal{A}$ is now 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:

$$T^* \geq \frac{\gamma^2}{9} \frac{\lambda A(1+\frac{1}{q})^2}{m^2 v^2 (1-\rho-\frac{2\lambda\bar{r}}{mqv})^2} - \frac{\bar{s}(1-2\rho)}{2\rho},$$

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 the same numerical constant from the uncapacitated bound. Note that for the case $q \to \infty$ and $m = 1$ this reduces to our earlier bound, albeit with a weaker constant.

When $m > 1$ and $q = \infty$, 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 a single-server policy.

For q finite, we construct policies, μ, for which

$$T_\mu \sim \gamma_\mu^2 \frac{\lambda A(1-\frac{1}{q})^2}{m^2 v^2 (1-\rho-\frac{2\lambda\bar{r}}{mqv})^2} \quad \text{as} \quad \rho + \frac{2\lambda\bar{r}}{mqv} \to 1,$$

and therefore have a constant factor guarantee. 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. A better guarantee is provided by a policy based on tour partitioning adapted from the static heuristic analyzed by Haimovich and Rinnooy Kan [20]. When there are k depots, these results also hold under certain symmetry conditions.

These results provide some intuitively satisfying insights. For example, when $m = 1$ and $q < \infty$ they imply a necessary and sufficient condition for stability is

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

Observe that this condition is no longer independent of the service region geometry because of the presence of $\overline{r}$; however, for $q \rightarrow \infty$ the dependence vanishes.

The second term in this stability condition has the interpretation of a *radial collection cost* in the sense of Haimovich and Rinnooy Kan [20]. That is, $2\overline{r}/v$ 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 travel time 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 the stability difference, $1 - \rho - \frac{2\lambda\overline{r}}{vq}$, just as it does in the uncapacitated case. Note that the average radial distance $\overline{r}$ plays a crucial role in the system's behavior in this case. Indeed, we prove that if one has the option of locating the depot anywhere within $\mathcal{A}$, then minimizing $\overline{r}$ (i.e. locating the depot at the median) is always optimal in heavy traffic.

1.4.3 General Distributions

All the heavy traffic results previewed thus far apply to the case of uniformly distributed demand (*i.e.* $f(x) = 1/A$) and Poisson arrivals. In Chapter 5, we take a different tact and investigate the m-server, uncapacitated problem where demands are distributed in the service regions $\mathcal{A}$ according to a general continuous density $f(x)$ and arrivals occur according to a general renewal process. Since the proofs for the lower bounds in Chapters 3 and 4 relied heavily on the Poisson and uniformity assumptions, we are forced to develop different proof techniques in this chapter.

In the general-$f(x)$ case, in turns out that we must distinguish between policies which are *spatially fair* (*i.e.* those providing the same mean system time for all locations within $\mathcal{A}$) and policies which are *discriminatory* (*i.e.* those providing different mean system times to different regions of $\mathcal{A}$). Within the class of spatially fair policies, we show that

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

where $\gamma \geq \frac{2}{3\sqrt{\pi}}$. As mentioned, this improves on the constant for the uniform case, for which $\left[\int_A f^{1/2}(x)dx\right]^2 = A$. Note, however, that this bound is an asymptotic bound while the previous bound is valid for all ρ. We then show that natural extensions of the TSP policies for the uniform case gives the same $\frac{\beta^2}{2\gamma^2}$ guarantee as in the uniform case.

For discriminatory policies, we show that

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

where γ is the same as in the fair case. We then propose a policy with a $\frac{\beta^2}{2\gamma^2}$ guarantee for the special case where $f(x)$ is piecewise uniform. We conclude Chapter 5 by discussing the relationship between the fair and discriminatory behaviors, giving some numerical results on the behavior or the space filling curve and nearest neighbor policies for general $f(x)$ and outlining extensions to the capacitated vehicle problem.

1.4.4 Extensions and Conclusions

Finally, in Chapter 6 we give some results and observations on a variety of extensions to the basic problem. We begin by discussing extensions to higher dimensions. We then look at the problem of minimizing a joint travel cost and system time objective and show that the basic tradeoff between the two is an inherent feature of most of our proposed policies. Then, we briefly discuss some stylized planning models based on our results. Lastly, we discuss using our lower bounds to analyze other combinatorial problems in the plane. In Chapter 7 we give our conclusions.

Chapter 2

Probabilistic and Queueing Background

In this chapter, we present some results from geometrical probability and queueing theory that are used frequently in the remaining chapters.

2.1 Jensen's Inequality

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

$$E[f(X)] \geq f(E[X]), \tag{2.1}$$

provided the expectations exist.

2.2 Wald's Equation

Let $\{X_i; i \geq 1\}$ be a sequence of i.i.d. random variables with $E[X] < \infty$ and N be a finite-mean random variable with the property that $P\{N = n\}$ is independent of $\{X_i; i > n\}$ for all n. (Such a random variable N is said to be a *stopping time* for the sequence $\{X_i; i \geq 1\}$.) Then

$$E\left[\sum_{i=1}^{N} X_i\right] = E[N]E[X_i]. \tag{2.2}$$

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

In a $GI/G/1$ queue let $\frac{1}{\lambda}$ be the expected interarrival time and $\overline{s}$ be the expected service time. Let σ_a^2 and σ_s^2 denote the variances of the interarrival and service time distribution respectively. The traffic intensity is $\rho = \lambda\overline{s}$. There is no simple explicit expression for the expected waiting time W in this case. (The average system time T is simply $W + \overline{s}$.) However, Kingman [26] (see also Kleinrock [29]) proves that

$$W \leq \frac{\lambda(\sigma_a^2 + \sigma_s^2)}{2(1-\rho)}. \tag{2.3}$$

In addition, this upper bound is asymptotically exact as $\rho \rightarrow 1$. For the $M/G/1$ it is well known (see Kleinrock [29]) that

$$W = \frac{\lambda\overline{s^2}}{2(1-\rho)}, \tag{2.4}$$

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

2.4 A Heavy Traffic Limit for the Waiting Time in a $GI/G/m$ Queue

Kingman [27] gives the following limit for the waiting time W in a G/G/m queue (*c.f.* [30]),

$$W \sim \frac{\lambda(\sigma_a^2 + \sigma_s^2/m^2)}{2(1-\rho)} \qquad \text{as } \rho \rightarrow 1, \tag{2.5}$$

where as above σ_a^2 and σ_s^2 are the variances for the interarrival times and service time respectively, $1/\lambda$ is the mean interarrival time, $\bar{s}$ is the mean service time and $\rho = \frac{\lambda\bar{s}}{m}$.

2.5 Symmetric Cyclic Queues

Consider a queueing system that consists of k queues $Q_1, Q_2, \ldots, Q_k$ each with infinite capacity. Customers arrive at each queue according to independent Poisson processes with the same arrival intensity λ/k. The queues are served by a single server who visits the queues in a fixed cyclic order $Q_1, Q_2, \ldots, Q_k, Q_1, Q_2, \ldots$. The travel time around the cycle is a constant d. The service times at every queue are independent, identically distributed random variables with mean $\overline{s}$ and second moment $\overline{s^2}$. The traffic intensity is $\rho = \lambda\overline{s}$. The server uses the exhaustive service policy, i.e. servicing each queue i until the queue is empty before proceeding. The expected waiting time for this system is given by (see Bertsekas and Gallager [8], p.156)

$$W = \frac{\lambda\ \overline{s^2}}{2(1-\rho)} + \frac{(1-\frac{\rho}{k})}{2(1-\rho)}d. \tag{2.6}$$

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 [16]).

2.6 A Heavy Traffic Limit for the $\sum$GI/G/m Queue

A queue is denoted $\sum$GI/G/m if its input process is the superposition of k independent renewal processes (not necessarily identical). The following theorem is due to Inglehart and Whitt [22] (*c.f.* Flores [17]):

Theorem 2.1 (Inglehart and Whitt [22]) *Consider an m server queue fed by the superposition of k renewal processes. Let $1/\lambda_i$ and $\sigma_{a_i}^2$ denote, respectively, the mean and variance of the interarrival time of the i-th renewal process, $i = 1, 2, \ldots, k$. Let $1/\mu_j$ and $\sigma_{b_j}^2$ denote the mean and variance, respectively, of the service times at server $j = 1, 2, \ldots, m$. Define $\lambda \equiv \sum_{i=1}^{k} \lambda_i$ $\mu \equiv \sum_{j=1}^{m} \mu_j$ and $\rho \equiv \frac{\lambda}{\mu}$. Then as $\rho \rightarrow 1$ the mean waiting time in queue, W, satisfies*

$$W \sim \frac{\sum_{i=1}^{k} \lambda_i^3 \sigma_{a_i}^2 + \sum_{j=1}^{m} \mu_j^3 \sigma_{b_j}^2}{2\mu^2(1-\rho)}. \tag{2.7}$$

2.7 Geometrical Probability

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

$$E[\|X_1 - X_2\|] = c_1\sqrt{A}, \quad E[\|X_1 - X_2\|^2] = c_2 A, \tag{2.8}$$

where $c_1 \approx 0.52, c_2 = \frac{1}{3}$ (see Larson and Odoni [31], p.135). If we let x^* denote the center of a square of area A, then it is known [31] that the first and second moment of the distance to a uniformly chosen point X are given by

$$E[|X - x^*|] = c_3\sqrt{A}, \quad E[|X - x^*|^2] = c_4 A, \tag{2.9}$$

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

2.8 Asymptotic Properties of the TSP in the Euclidean Plane

Let $X_1 \dots X_n$ 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 the following theorem holds:

Theorem 2.2 (Bearwood, Halton and Hammersley [6]) *There exists a constant, β, such that*

$$\lim_{n\to\infty} \frac{L_n}{\sqrt{n}} = \beta\sqrt{A}. \tag{2.10}$$

with probability one. (See also [44] and [34].)

Suppose now that $X_1 \dots X_n$ are distributed according to a general distribution that has compact support and an absolutely continuous part $f(x)$, then

$$\lim_{n\to\infty} \frac{L_n}{\sqrt{n}} = \beta \int f^{1/2}(x) \tag{2.11}$$

(See [44].)

In his recent experimental work with very large scale TSP's, Johnson [24] estimated $\beta_{TSP} \approx 0.72$. Also, for finite n there exists a constant $\bar{\beta}$ such that with probability one

$$L_n \leq \bar{\beta}\sqrt{An}. \tag{2.12}$$

This can be shown using, for example, the strip heuristic in which case $\bar{\beta} = 2$ [25]. In addition, it is also well known (see [34], p. 189) that $\lim_{n\to\infty} Var(L_n) = O(1)$, and therefore

$$\lim_{n\to\infty} \frac{Var(L_n)}{n} = 0. \tag{2.13}$$

2.9 Space Filling Curves

The following results are due to Platzman and Bartholdi [37]. Let $\mathcal{C} = \{\theta | 0 \leq \theta \leq 1\}$ denote the unit circle and $\mathcal{S} = \{(x, y) | 0 \leq x \leq 1, 0 \leq y \leq 1\}$ denote the unit square. Then there exists a continuous mapping ψ from $\mathcal{C}$ onto $\mathcal{S}$ with the property that for any $\theta, \theta' \in \mathcal{C}$,

$$||\psi(\theta) - \psi(\theta')|| \leq 2\sqrt{|\theta - \theta'|}. \tag{2.14}$$

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

$$L_n \leq 2\sqrt{n}. \tag{2.15}$$

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

$$\limsup_{n\to\infty} \frac{L_n}{\sqrt{n}} = \beta_{SFC} \tag{2.16}$$

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

Chapter 3

The Single Uncapacitated Vehicle DTRP

In this chapter we examine the simplest case of the DTRP in which demands are uniformly distributed in $\mathcal{A}$ and the entire region is served by a single vehicle, and no capacity constraints are imposed. In §3.2 we derive lower bounds on the system time. Then in §3.3 we propose some specific policies and compare their performance to these lower bounds in light and heavy traffic. A numerical example is given in §3.5 to illustrate the relative performance of the proposed policies.

3.1 Notation

Locations in $\Re^2$ are generically denoted x (*i.e.* $x = [x_1 \; x_2]$) and random locations are denoted X. Subsets of $\Re^2$ are typically denoted by calligraphic letters (*e.g.* $\mathcal{S}$). 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. This association is in a sense arbitrary. Indeed, in Chapter 5 we shall reverse this convention and define d_i to be the distance traveled from demand i to demand $i+1$. In the limit as $i \longrightarrow \infty$ these two definition are equivalent.

The limiting expected values of these random variables are defined by $\overline{s} = \lim_{i\to\infty} E[s_i]$, $W = \lim_{i\to\infty} E[W_i]$, $T = \lim_{i\to\infty} E[T_i]$ and $\overline{d} = \lim_{i\to\infty} E[d_i]$. We shall assume that if the service policy is stable (That is, the number in the system is bounded almost surely for all times t.), then the system is ergodic and these limits exist. Thus, for example, we can alternatively interpret $\bar{d}$ as the limit as $t \to \infty$ of the cumulative distance traveled in $[0,t]$ divided by the number of demands served in $[0,t]$.

For a given policy μ, we write T_μ to denote $\lim_{i\to\infty} E[T_i]$ under this policy. The optimal system time within the class $\mathcal{M}$ of stable policies is denoted T^*. That is,

$$T^* = \min_{\mu\in\mathcal{M}} T_\mu.$$

3.2 Lower Bounds on the Optimal DTRP Policy

We first establish two simple but powerful lower bounds on the optimal expected system time, T^*.

3.2.1 A Light Traffic Lower Bound

The following bound is most useful in the case of light traffic ($\lambda \to 0$):

Theorem 3.1

$$T^* \geq \frac{E[||X - x^*||]}{v(1-\rho)} + \frac{\lambda \overline{s^2}}{2(1-\rho)} + \overline{s}, \tag{3.1}$$

where x^* is the median of the region $\mathcal{A}$.

We note for the special case where $\mathcal{A}$ is a square,

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

Proof

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

$$T = W^d + W^s + \overline{s}, \tag{3.2}$$

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

To bound W^d, note that it is at least as large as the travel delay between the server's location at the time of a demand's arrival and the demand'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 below by the expected delay between a server location selected from this distribution and a uniform location. Now, suppose we had the option of locating the server in the best *a priori* location, x^*; that is, the location that minimizes the expected travel time to a uniformly chosen location, X. This certainly yields a lower bound on the expected distance between the server and the arrival, so

$$W^d \geq \frac{1}{v} \min_{x \in \mathcal{A}} E[\|X - x\|]. \tag{3.3}$$

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 (2.9),

$$W^d \geq \frac{c_3}{v}\sqrt{A} \approx 0.383\sqrt{A}. \tag{3.4}$$

To bound W^s, let N denote the expected number of demands served during a waiting time. Since service times are independent, we then have

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

where the second term is the expected residual service time of the demand being served at the time of arrival. Since in steady state the expected number of demands served during a wait is equal to the expected number that arrive, we can apply Little's law to get

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

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

$$W^s = \frac{\rho}{1-\rho}(W^d) + \frac{\lambda\overline{s^2}}{2(1-\rho)}. \tag{3.5}$$

Combining (3.2), (3.3) and (3.5) and noting that these bounds are true for *all* policies we obtain Theorem 3.1.

□ (Theorem 3.1)

We note that uniformity is only used to evaluate $E[||X - x||]$ for square regions, and thus Theorem 3.1 holds for $X \sim f(x)$ as well.

3.2.2 A Heavy Traffic Lower Bound

A lower bound that is most useful for the heavy traffic ($\rho \to 1$) case is provided by the following theorem:

Theorem 3.2 *There exists a constant γ such that*

$$T^* \geq \gamma^2 \frac{\lambda A}{v^2(1-\rho)^2} - \frac{\bar{s}(1-2\rho)}{2\rho}. \tag{3.6}$$

Proof

We begin with the following important lemma:

Lemma 3.1

$$\overline{d} \geq \gamma \frac{\sqrt{A}}{\sqrt{N + 1/2}}, \tag{3.7}$$

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

Before proving this lemma, observe that Theorem 3.2 is easily derived from it by substituting the bound on $\overline{d}$ into the stability condition,

$$\overline{s} + \frac{\overline{d}}{v} \leq \frac{1}{\lambda}, \tag{3.8}$$

which yields

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

After rearranging, noting that $T = W + \overline{s}$ and $N = \lambda W$, we obtain the bound of Theorem 3.2. Thus, Theorem 3.2 is established once Lemma 3.1 is proven.

□ (Theorem 3.2)

Proof of Lemma 3.1

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.*

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

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

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

$$P\{Z_i \leq z\} \leq \frac{\pi z^2}{A}, \tag{3.9}$$

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 $Z^* \equiv \min\{Z_0^*, Z_1, ..., Z_{N_1}\}$. Hence,

$$\overline{d} \geq E[Z^*]. \tag{3.10}$$

We next bound the right hand side of (3.10). 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

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

Since $E[N_1] = N$ and $E[I_{Z_i}] = P\{Z_i \leq z\}$ is bcunded according to (3.9), we obtain

$$P\{Z^* > z\} \geq 1 - P\{Z_0^* \leq z\} - N\frac{\pi z^2}{A}. \tag{3.11}$$

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

Lemma 3.2 : $P\{Z_0^* \leq z\} \leq \frac{\pi z^2}{A}(N+1)$.

Proof

First, consider *any* set $\mathcal{S}$ of n points in $\mathcal{A}$. Let X be a uniformly distributed location in $\mathcal{A}$ independent of $\mathcal{S}$ and define $Z^* \equiv \min_{x \in \mathcal{S}} \|X - x\|$. For each point in $\mathcal{S}$, construct a circle of radius z centered at the point, and let $A(\mathcal{S})$ denote the total area in $\mathcal{A}$ covered by the intersection of these circles. Then,

$$P\{Z^* \leq z\} = \frac{A(\mathcal{S})}{A} \leq \frac{\pi z^2}{A} n.$$

Since X_0 is independent of $\mathcal{S}_0$ under any condition on $\mathcal{S}_0$, we can condition on the value N_0 and use the the above bound to assert that

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

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

□ (Lemma 3.2)

Using the result of Lemma 3.2 in (3.11) yields

$$P\{Z^* > z\} \geq 1 - \frac{\pi}{A}(2N + 1)z^2.$$

Combining this with the trivial bound $P\{Z^* > z\} \geq 0$ we obtain

$$E[Z^*] \geq \int_0^\infty \max\{0, 1 - \frac{\pi}{A}(2N + 1)z^2\}dz = \int_0^{\sqrt{1/c}} (1 - cz^2)dz,$$

where $c \equiv \frac{\pi(2N+m)}{A}$. The integral gives $\frac{2}{3}c^{-1/2}$, whereupon substituting the value c we establish Lemma 3.1 with $\gamma = \frac{2}{3\sqrt{2\pi}} \approx 0.266$.

□ (Lemma 3.1)

A few comments on the lower bound of Theorem 3.2 are in order. First, it shows that the waiting time grows at least as fast as $(1-\rho)^{-2}$ rather than $(1-\rho)^{-1}$ as is the case for a classical queueing system. Also, it is only a function of the first moment of the on-site service time, which is again a significant departure from traditional queueing system behavior (e.g. the $M/G/1$ system).

The explanation for this behavior lies in the geometry of the system. The bound of Lemma 3.1 gives the minimum average number of demands, N, that must be maintained in the system to reduce $\bar{d}$, the average travel distance per demand, to a given value. Note that this bound on $\bar{d}$ is inversely proportional to the *square root* of N. At the same time, the stability condition given by Equation (3.8) requires that

$$\bar{d} \leq \frac{v(1-\rho)}{\lambda}.$$

Thus, N must grow like $\frac{1}{(1-\rho)^2}$ to maintain stability, which is a much more rapid than the increase in N due simply to traditional queueing-type delays.

Because several loose assumptions were used in the above proof (*e.g.* bounding probabilities by expectations), it is likely that the value $\gamma \approx 0.266$ is not tight. For example, if one assumes locations of demands at service completion epochs are approximately uniform, then by a modified argument one can obtain $\gamma = 1/2$. Indeed, in Chapter 5 we develop a different proof of Lemma 3.1 for general demand distributions that improves the constant, raising it to $\gamma \approx 0.376$. This bound, however, is asymptotic in the sense that it only applies for $N \to \infty$, while the bound in Lemma 3.1 applies for all N.

3.3 Some Proposed Policies for the DTRP

In this section, we propose and analyze several policies for the single vehicle DTRP. The first class of policies is 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 partitioning policy based on subdividing the large square $\mathcal{A}$ into smaller squares, each of which is served locally using a FCFS discipline. Using results on cyclic queues, we show that the this policy is within a constant factor of the lower bounds for all values of $\rho < 1$. This also establishes $\rho < 1$ as a sufficient (as well as, obviously, necessary) condition for stability in the sense that there exist stable policies for every $\rho < 1$.

We next introduce a more sophisticated policy based on forming successive TSP tours. Its average system time is nearly half that of the partitioning policy. A further improvement on this policy reduces the system time by another factor of two, and this gives us our best policy. Next, a policy based on space filling curves is examined. It too has a constant factor performance guarantee and is shown via simulation to have a system time about 70% greater than the best TSP policy. Finally, we examine the policy of serving the nearest neighbor. Because of analytical difficulties, we simulate it and show the average system time is about 60% greater than the best TSP policy. Despite their relatively poor performance guarantees, these policies are computationally efficient and easy to implement. Thus, they may be attractive in practice.

3.3.1 FCFS Policies

The simplest policy for the DTRP is to service demands in the order in which they arrive (FCFS). The first policy we examine of this type we only analyze for $\mathcal{A}$ a square. It is defined as follows:

The FCFS Policy

When demands are present, travel directly from one demand location to the next following a FCFS order. When no unserved demands are present following a service completion, wait at the current location until a new demand arrives before beginning restarting FCFS service.

The system time of this policy, T_{FCFS}, is described by the following proposition:

Proposition 3.1 *If $\mathcal{A}$ is a square, then*

$$\frac{T_{FCFS}}{T^*} \leq \frac{c_1}{c_3} \approx 1.36 \qquad \text{as } \lambda \to 0.$$

where c_1 and c_3 are defined as in Equations (2.8) and (2.9) respectively.

Proof

Because demand locations are independent of the order of arrivals and also the number of demands 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, because each d_i is simply the distance between two independent, uniformly distributed locations in $\mathcal{A}$. Therefore, the Pollaczek-Khinchin (P-K) formula (2.4) still holds. (See [8] page 142-143 for a proof of the P-K formula that does not require mutual independence of service times.)

The first and second moments of the total service requirement are, by (2.8), $\bar{s} + c_1\sqrt{A}/v$ and $\overline{s^2} + 2c_1\sqrt{A}\bar{s}/v + c_2 A/v^2$ respectively, where $c_1 \approx 0.52, c_2 = \frac{1}{3}$. The average system time is therefore, by the P-K formula (2.4),

$$T_{FCFS} = \frac{\lambda(\overline{s^2} + 2c_1\sqrt{A}\bar{s}/v + c_2 A/v^2)}{2(1 - \lambda c_1\sqrt{A}/v - \rho)} + \bar{s} + c_1\sqrt{A}/v. \qquad (3.12)$$

The stability condition for this policy is $\rho + \lambda c_1\sqrt{A}/v < 1$; therefore, this policy is unstable for values of ρ approaching 1. For $\lambda \rightarrow 0$, the first term in (3.12) approaches zero. Likewise, the second term in the bound of Theorem 3.1 also approaches zero as $\lambda \rightarrow 0$. So for the light traffic case we have

$$\frac{T_{FCFS}}{T^*} \leq \frac{\overline{s} + c_1\sqrt{A}/v}{\overline{s} + c_3\sqrt{A}/v}, \qquad \text{as } \lambda \rightarrow 0.$$

Since $\overline{s}$ could be arbitrarily small, the worst case relative performance for this policy in light traffic is as described by the proposition.

□ (Proposition 3.1)

The FCFS policy can be modified and generalized to yield asymptotically optimal performance in light traffic as follows:

The Stochastic Queue Median (SQM) Policy

Locate the server at x^*, the median of $\mathcal{A}$. That is, x^* is the location x that minimizes $E[||X - x||]$. Service each demand FCFS by traveling directly to the demand's location from the median, servicing the demand, and then returning back to the median after service is completed. If no demands are present in the system, the server waits at the median until the next demand arrives.

Proposition 3.2

$$\frac{T_{SQM}}{T^*} = 1 \qquad \text{as } \lambda \rightarrow 0.$$

Proof

We first prove it for the case where $\mathcal{A}$ is square. 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 time between the median and the service location. The system time of an individual demand, however, includes the wait in queue plus the *one way* travel time to the service location plus the on-site service. Therefore, the average system time under this policy is given by (*c.f.* Equation 2.9)

$$T_{SQM} = \frac{\lambda(\overline{s^2} + 4c_3\sqrt{A}\overline{s}/v + 4c_4A/v^2)}{2(1 - 2\lambda c_3\sqrt{A}/v - \rho)} + \overline{s} + c_3\sqrt{A}/v, \tag{3.13}$$

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

Letting λ approach zero, the first term above goes to zero, and since c_3 is the constant of the lower bound in Theorem 3.1 we get

$$\frac{T_{SQM}}{T^*} \longrightarrow 1, \qquad as\ \lambda \longrightarrow 0. \tag{3.14}$$

This argument can be generalized to arbitrary regions $\mathcal{A}$ by substituting $E[||X - x^*||]$ for $c_3\sqrt{A}$ and $E[||X - x^*||^2]$ for c_4A in (3.13).

□ (Proposition 3.2)

This is an intuitively satisfying if not altogether surprising result. It is conjectured by Psaraftis in [39]. It is also analogous to the results achieved by Berman et. al. [7] and Batta et. al. [5] for the optimal location of a server on a network operated under a FCFS policy. Our result is somewhat stronger because our lower bound is on *all* policies, not just FCFS policies. Therefore it establishes not only the optimality of the median location for the SQM, but also the optimality of the SQM discipline itself.

The FCFS and SQM policies become unstable for $\rho \longrightarrow 1$. The reason for this is that the average distance traveled per service, $\overline{d}$, remains fixed, yet the stability condition (3.8) implies $\overline{d} \leq \frac{v(1-\rho)}{\lambda}$, so $\overline{d}$ must decrease as ρ (and λ) are increased. As shown below, a policy that is stable for all values of ρ must increasingly restrict the distance the server is willing to travel between services as the traffic intensity increases.

3.3.2 The Partitioning Policy

In this section we examine a policy that achieves the restriction on $\overline{d}$ mentioned above through a partition of the service region $\mathcal{A}$. The analysis relies on results

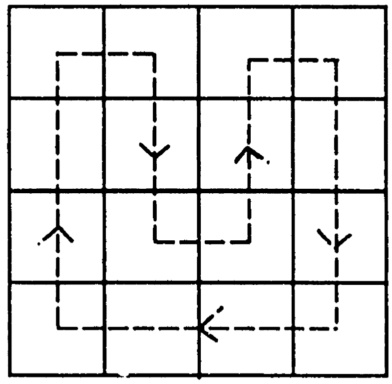

Figure 3.1: Sequence for Serving Subregions PART Policy ($m = 4$)

for symmetric, cyclic queues, so readers unfamiliar with this area are encouraged to reexamine the definitions and results in Chapter 2.

Suppose $\mathcal{A}$ is a square. Consider the following policy:

The Partition (PART) Policy

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, serve demands using a FCFS discipline. The server services a subregion until there are no more demands left in that subregion. It then moves on to the next subregion and services it until no more demands are left, etc. The sequence of regions the server follows is shown in Figure 3.1 for the case $m = 4$. (Note that the server always moves to an *adjacent* subregion.) The pattern is continuously repeated.

Proposition 3.3

$$\frac{T_{PART}}{T^*} \leq \frac{2c_1}{\gamma^2} \quad as\ \rho \rightarrow 1$$

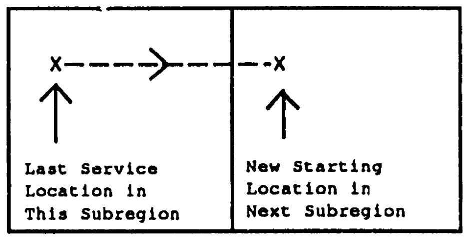

Figure 3.2: PART Projection Policy for Moving to Adjacent Subregion

Proof

We must clarify how the server moves from one subregion to the next. We assume the server uses the projection rule shown in Figure 3.2. Its last location in a given subregion is simply "projected" onto the next subregion to determine the server's new starting location. The server then travels in a straight line between these two locations. As a result of this rule, note that the distance traveled between subregions is a constant $\frac{\sqrt{A}}{m}$, and that each starting location is uniformly distributed and independent of the locations of demands in the new subregion. These properties of the starting location simplify the analysis. In practice, one might use a more intelligent rule such as moving directly to the first demand in the new subregion. The total travel distance of this tour is $m^2(\sqrt{A}/m) = m\sqrt{A}$.

Notice that to construct the pattern shown in Figure 3.1, m must be even. If m is odd, the server ends up in the upper right subregion and must travel to the lower right subregion to restart the cycle. This adds an additional $\sqrt{A} - \sqrt{A}/m$ to the total travel distance. To simplify the analysis, we use only the expression for even m. As shown below, m must be large in heavy traffic, so for $\rho \rightarrow 1$ the relative 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 $\overline{s} + c_1\frac{\sqrt{A}}{vm}$ and $\overline{s^2} + 2c_1\overline{s}\frac{\sqrt{A}}{vm} + c_2\frac{A}{v^2m^2}$ 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 $m\sqrt{A}/v$ 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 [8] still holds. Recalling that the expression in (2.6) is for the waiting time in queue only, the average system time for this policy is given by

$$T_{PART} = \frac{\lambda(\overline{s^2} + 2c_1\overline{s}\frac{\sqrt{A}}{vm} + c_2\frac{A}{v^2m^2})}{2(1-\lambda(\overline{s}+c_1\frac{\sqrt{A}}{vm}))} + \frac{1-\frac{\lambda}{m^2}(\overline{s}+c_1\frac{\sqrt{A}}{vm})}{2(1-\lambda(\overline{s}+c_1\frac{\sqrt{A}}{vm}))} m\frac{\sqrt{A}}{v} + c_1\frac{\sqrt{A}}{vm} + \overline{s}. \quad (3.15)$$

The stability condition is

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

Defining the critical value m_c by

$$m_c \equiv \frac{c_1\lambda\sqrt{A}}{v(1-\rho)}, \quad (3.16)$$

the stability conditions becomes $m > m_c$. Note that for any $\rho < 1$ we can find an $m > m_c$ such that this policy is stable. Theorem 3.2 shows that $\rho < 1$ is a necessary condition for stability. Thus, since the optimal policy has a waiting time no greater than the PART policy, we have the following theorem:

Theorem 3.3 *There exists a stable optimal policy for the single, uncapacitated vehicle DTRP if and only if* $\rho < 1$.

Since ρ is determined only by the on-site service mean and the arrival rate, we see that the service region characteristics (size, shape, etc.) do not affect the amount of traffic the system can support (provided the service region is bounded, of course). This condition is perhaps surprising at first, since one might assume that the geometry should play a role here. Intuitively, it says that as long as the vehicle has some slack time left after tending to on-site service, then one can find a stable policy that reduces the travel time so as to fit within this slack time. Though the geometry does not determine stability, it certainly affects the resulting system time as we shall see next.

For given system parameters λ, $\overline{s}$, $\overline{s^2}$ and A, one could perform a one dimensional optimization over $m > 1$ using (3.15) to get the optimum number of partitions; however, since equation (3.15) is quite complicated, we concentrate on finding the optimal value, m^*, for the heavy traffic case.

For ($\rho \rightarrow 1$), (3.16) implies that any feasible m is large ($m > m_c$). Therefore ignoring the $O(1/m)$ and smaller terms in the numerators of (3.15) we obtain

$$T_{PART} \approx \frac{\lambda\overline{s^2} + m\frac{\sqrt{A}}{v}}{2(1-\rho-\lambda c_1\frac{\sqrt{A}}{vm})} = \frac{m^2\frac{\sqrt{A}}{v} + m\lambda\overline{s^2}}{2(m(1-\rho)-\lambda c_1\frac{\sqrt{A}}{v})}. \tag{3.17}$$

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

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

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

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

If we substitute this value into (3.17), then in heavy traffic

$$T_{PART} \approx 2c_1\frac{\lambda A}{v^2(1-\rho)^2} + \frac{\lambda\overline{s^2}}{1-\rho}. \tag{3.18}$$

For $\rho \rightarrow 1$, the first term above dominates. Comparing to the bound in Theorem 3.2) establishes the proposition.

□ (Proposition 3.3)

Proposition 3.3 says that T_{PART} is within a constant factor of the optimum in heavy traffic, though the provable factor is indeed quite large (about 7 using $\gamma \approx 0.376$). Note also that the geometrical constant c_1 and the area A certainly affect the system time of this policy.

3.3.3 The Traveling Salesman Policy

The traveling salesman policy (TSP for short) is based on collecting demands into sets that can then be served using an optimal TSP tour. It is defined as follows:

The Traveling Salesman (TSP) Policy

Let $\mathcal{N}_k$ denote the *kth* set of n demands to arrive, where n is a given constant that parameterizes the policy, e.g. $\mathcal{N}_1$ is the set of demands $1,\ldots,n$, $\mathcal{N}_2$ is the set of demands $n+1,\ldots,2n$, etc. Assume the server operates out of a depot at a random location in $\mathcal{A}$. When all demands in $\mathcal{N}_1$ have arrived, we form an optimal (shortest length) tour on these demands starting and ending at the depot. Service demands by following the tour. If all demands in $\mathcal{N}_2$ have arrived when the tour of $\mathcal{N}_1$ is completed, these are serviced using a TSP tour; otherwise, the server waits until all $\mathcal{N}_2$ demands arrive before serving it. In this manner, service sets in a FCFS order. Optimize over n.

Proposition 3.4

$$\frac{T_{TSP}}{T^*} \leq \frac{\beta^2}{\gamma^2} \quad \text{as } \rho \to 1,$$

where $\beta \approx 0.72$ is the Euclidean TSP constant.

Proof

Suppose one considers the set $\mathcal{N}_k$ to be the *kth* "customer". Since the interarrival time (time for n new demands to arrive) and service time (n on-site 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 and the n on-site service times. If we let L_n denote the length of such a tour and $E[L_n]$ 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 $\frac{E[L_n]}{v} + n\bar{s}$ and the variance is $\frac{Var[L_n]}{v^2} + n\sigma_s^2$, where $\sigma_s^2 = \overline{s^2} - \bar{s}^2$ is the variance of the on-site service time.

We are now in a position to apply the $GI/G/1$ upper bound (see Equation (2.3)) for the average waiting time of sets, W_{set}. This giv s

$$W_{set} \leq \frac{\frac{\lambda}{n}(\frac{n}{\lambda^2} + \frac{Var[L_n]}{v^2} + n\sigma_s^2)}{2(1 - \frac{\lambda}{n}(\frac{E[L_n]}{v} + n\bar{s}))} \tag{3.19}$$

$$= \frac{\lambda(1/\lambda^2 + \frac{Var[L_n]}{v^2 n} + \sigma_s^2)}{2(1 - \rho - \lambda\frac{E[L_n]}{vn})}. \tag{3.20}$$

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 (2.10) and (2.13) that as $n \to \infty$

$$\frac{E[L_n]}{n} \sim \beta\frac{\sqrt{A}}{\sqrt{n}}, \tag{3.21}$$

and

$$\frac{Var[L_n]}{n} \sim 0. \tag{3.22}$$

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

$$W_{set} \leq \frac{\lambda(1/\lambda^2 + \sigma_s^2)}{2(1 - \rho - \lambda\beta\frac{\sqrt{A}}{v\sqrt{n}})}. \tag{3.23}$$

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

$$n > \frac{\lambda^2\beta^2 A}{v^2(1-\rho)^2}. \tag{3.24}$$

For $\rho \to 1$, the above implies that n *must* be large, and thus our use of asymptotic TSP results is indeed justified.

The waiting time given in (3.23) is not itself an upper bound on the wait for service of an individual demand; it is the wait in queue for a set. The time of arrival of a set is actually the time of arrival of the last demand in that set. Therefore, we must add to (3.23) the time a demand waits for its set to form, denoted W^-, and also the time it takes to complete service of the demand once its set enters service, denoted W^+. By conditioning on the position that a given demand takes within its set, it is easy to show that

$$W^- = \frac{n-1}{2\lambda} \leq \frac{n}{2\lambda}.$$

By doing the same conditioning and noting that the travel time around the tour is no more that the length of the tour itself, we obtain

$$\begin{aligned} W^+ &\leq \frac{1}{n}\sum_{k=1}^{N} k\bar{s} + \bar{\beta}\frac{\sqrt{nA}}{v} \\ &\leq \frac{n}{2}\bar{s} + \bar{\beta}\frac{\sqrt{nA}}{v}, \end{aligned}$$

where $\bar{beta}$ is a bound on the TSP constant for finite n as in (2.12). Therefore, if the total system time is denoted T_{TSP},

$$T_{TSP} \leq \frac{\lambda(1/\lambda^2 + \sigma_s^2)}{2(1 - \rho - \lambda\beta\frac{\sqrt{A}}{v\sqrt{n}})} + \frac{n(1+\rho)}{2\lambda} + \bar{\beta}\frac{\sqrt{nA}}{v}. \tag{3.25}$$

We would like to minimize (3.25) with respect to n to get the least upper bound. (One can verify that (3.25) is convex for $n > 0$, so there is a unique minimum.) First, however, consider a change of variable

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

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

$$T_{TSP} \leq \frac{\lambda(1/\lambda^2 + \sigma_s^2)}{2(1-\rho)(1-y)} + \frac{\lambda\beta^2 A(1+\rho)}{2v^2(1-\rho)^2 y^2} + \frac{\lambda\beta\bar{\beta}A}{v^2(1-\rho)y}. \tag{3.26}$$

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

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

Substituting this approximation into (3.26) and noting that for $\rho \to 1$ the approximate y^* approaches one we have

$$T_{TSP} \leq \beta^2\frac{\lambda A}{v^2(1-\rho)^2} + \frac{\beta\lambda\sqrt{A(1/\lambda^2 + \sigma_s^2)}}{v(1-\rho)^{3/2}} + \frac{\beta\bar{\beta}\lambda A}{v^2(1-\rho)} \quad \rho \to 1.$$

Again, the leading term is proportional to $\frac{\lambda A}{v^2(1-\rho)^2}$. Therefore, comparing to Theorem 3.2 the proposition is established.

□ (Proposition 3.4)

The best estimate to date of β is approximately 0.72 [24], so the TSP policy has a system time in heavy traffic about half that of the partitioning policy. (In practice, heuristic rather than optimal tours would be used to reduce the computational burden, which would produce slightly higher system times.) These results suggest that the policy of forming successive TSP tours, which is a reasonable policy in practice, is also quite good theoretically. In addition to providing a theoretical guarantee, our analysis gives some guidance into optimally sizing routes for such policies by either minimizing the right hand side of (3.25) or using the approximate value for y^* given above.

3.3.4 The Modified TSP Policy

The following simple modification of the TSP policy can reduce the asymptotic system time by a factor of two:

The Modified TSP (MOD TSP) Policy

> Let k be a fixed positive integer. From a central point in the interior of $\mathcal{A}$, subdivide the service region into k wedges of area $1/k$. Within each subregion, form sets of size n/k. (n is a parameter to be determined.) As sets are formed, deposit them in a queue and service them FCFS with by forming a TSP on the set as in the TSP policy. Optimize over n.

The performance of this policy is given by the following proposition.

Proposition 3.5

$$\frac{T_{MOD\ TSP}}{T^*} \leq \frac{\beta^2}{2\gamma^2} \quad as\ \rho \rightarrow 1.$$

Proof

The analysis of this policy closely follows that of the TSP policy so we shall only outline the proof. To determine the waiting time in queue for a set, W_{set}, we note that each subregion generates sets that arrive according to independent renewal processes with rate $\hat{\lambda}_i = \frac{\lambda}{k}\frac{1}{(n/k)} = \frac{\lambda}{n}$ and interarrival time variance $\sigma_i^2 = (\frac{1}{(\lambda/k)})^2(\frac{n}{k}) = \frac{kn}{\lambda^2}$. Thus, the resulting queue is $\sum$GI/G/1 and we can make use of Theorem 2.1. Note that

$$\sum_{i=1}^{k} \hat{\lambda}_i^2 \sigma_i^2 = \frac{1}{\lambda}\hat{\lambda}^2$$

where $\hat{\lambda} \equiv \sum_{i=1}^{k} \hat{\lambda}_i = \frac{k\lambda}{n}$ is the overall arrival rate of sets. Thus, letting the random variable τ denote the time to service a set, Theorem 2.1 implies

$$W_{set} \sim \frac{\frac{1}{\lambda}\hat{\lambda}^2 + E^{-3}[\tau]\sigma_\tau^2}{2E^{-2}[\tau](1 - \hat{\lambda}E[\tau])} \quad \text{as } \hat{\lambda}E[\tau] \rightarrow 1.$$

Using the fact that $\hat{\lambda} \approx \frac{1}{E[\tau]}$ and $\hat{\lambda} = \frac{k\lambda}{n}$ this limit can be written

$$W_{set} \sim \frac{\lambda(\frac{1}{\lambda^2} + \frac{\sigma_\tau^2}{(n/k)})}{2(1 - \lambda\frac{E[\tau]}{(n/k)})} \quad \text{as } \lambda\frac{E[\tau]}{(n/k)} \rightarrow 1.$$

As before $\frac{\sigma_\tau^2}{(n/k)} \sim \sigma_s^2$ and $\frac{E[\tau]}{(n/k)} \sim \bar{s} + \beta\frac{\sqrt{A/k}}{v\sqrt{n/k}} = \bar{s} + \beta\frac{\sqrt{A}}{v\sqrt{n}}$. Thus, we obtain the same expression as in Equation (3.23) for W_{set}. (Though in this case it is a limit rather than an upper bound.) This shows the waiting time in queue as a function of n is unchanged asymptotically by the modification.

We next determine the effect of the modification on the other components of the waiting time. Note that though we have reduced the size of sets by $1/k$, the interarrival time in each region is a factor of k greater (*i.e.* the interarrival time mean is now k/λ). Thus, W^-, the wait for a set to form, is still bounded above by $\frac{n}{2\lambda}$. The wait for service once a set enters service, W^+, however, is now at most $\frac{n}{2k}\bar{s} + O(\sqrt{n})$ since set sizes have been reduced by $1/k$. Combining these terms and repeating the analysis we find that

$$T_{MOD\ TSP} \leq \beta^2\frac{\lambda A(1 + \frac{1}{k})}{2v^2(1 - \rho)^2} + \frac{\lambda\beta\sqrt{A(1/\lambda^2 + \sigma_s^2)}}{v(1 - \rho)^{3/2}} + O(\frac{1}{(1 - \rho)}) \quad \rho \rightarrow 1.$$

The leading term can be made arbitrarily close to $\frac{\beta^2}{2}\frac{\lambda A}{v^2(1-\rho)^2}$ by choosing k large and so by comparing to Theorem 3.2 the proposition is established.

□ (Proposition 3.5)

3.3.5 The Space Filling Curve Policy

We next analyze a policy based on space filling curves (the SFC policy). It was first proposed by Bartholdi and Platzman in [4]. The reader is encouraged to reexamine Chapter 2 for notation and basic results related to space filling curves. Let $\mathcal{C}$ and ψ be defined as in Chapter 2, and let the DTRP service region, $\mathcal{A}$, be a square of area A. Suppose we maintain the preimages of all demands in the system (i.e. their positions in $\mathcal{C}$). Then the SFC policy is to service demands as they are encountered in repeated clockwise sweeps of the circle $\mathcal{C}$. (Note that one could treat a depot as a permanent "demand" and visit it once per sweep.)

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

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

Taking expectation on both sides gives

$$W = E[N_0]\bar{s} + \frac{E[L]}{v} + \frac{\lambda \overline{s^2}}{2}. \tag{3.27}$$

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

$$E[L] \quad \leq \quad 2E[\sqrt{(N_0 + 2)A}] \tag{3.28}$$

$$\leq 2\sqrt{(N+2)A} \quad (Jensen's Ineq.)$$
$$\leq 2\sqrt{\lambda WA} + 2\sqrt{2A}.$$

Substituting these results into (3.27) we obtain the following quadratic inequality:

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

Solving the above for W and recalling that $T = W + \bar{s}$ we obtain

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

where $\gamma_{SFC} \leq 2$ and $o((1-\rho)^{-2})$ denotes terms that increase more slowly than $(1-\rho)^{-2}$ as $\rho \to 1$. Thus, by comparing to Theorem 3.2 we have the following proposition:

Proposition 3.6

$$\frac{T_{SFC}}{T^*} \leq \frac{\gamma_{SFC}^2}{\gamma^2} \leq \frac{4}{\gamma^2} \qquad as\ \rho \to 1.$$

This shows the SPC policy has a constant factor guarantee. The constant $\gamma_{SFC} = 2$ obtained by the above argument, however, is based on worst-case tours and is probably too large. If one assumes that the clockwise interval between the preimages of the server and the tagged demand is a uniform $[0,1]$ random variable and that the $\mathcal{N}_0$ points are approximately uniformly distributed on this interval, then a constant of $\gamma_{SFC} \approx \frac{2}{3}\beta_{SPC} \approx 0.64$ is obtained.

To estimate γ_{SFC} more precisely, we performed simulation experiments. The method of ***batch means*** (see [32]) was used to estimate the steady state value of T_{SFC}. In this method, demands are grouped into batches of a fixed size. If the batch size is large enough, the sample means from each batch are approximately uncorrelated and normally distributed [33]. (We used 200 times the minimum average number in the system given by Theorem 3.2 as our batch size.) The sample mean and variance of the individual batch means were then used in a t-test to estimate T_{SFC}. The simulation was terminated when the 99% confidence interval about the estimate

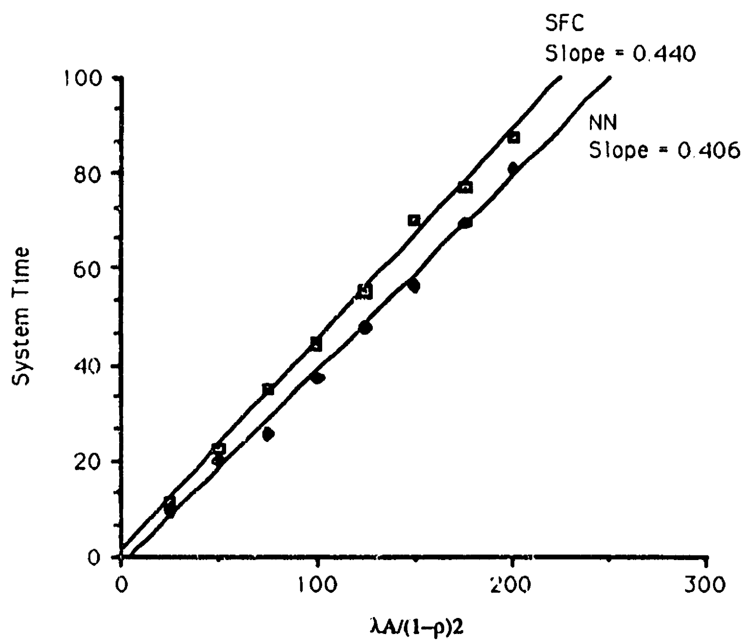

Figure 3.3: Simulation Results: T_{SFC} and T_{NN} vs. $\lambda A/(1-\rho)^2$

reached a width less than 10% of the value of the estimate. This method was selected because the busy periods of the SFC policy were quite long (indeed, almost nonterminating) at high utilization values, which precluded the use of techniques based on regeneration points.

The simulation was run for $A = 1$, $v = 1$ and a range of parameter values ρ, $\overline{s}$ and $\overline{s^2}$. Figure 3.3 shows one example of the simulation estimate of T_{SFC} plotted against $\lambda A/(1-\rho)^2$ for the case $\overline{s} = 0.1$ and $\overline{s^2} = 0.01$ (zero variance). Each point is a different value of ρ in the range $0.5 - 0.8$ The results showed that γ_{SFC} is approximately 0.66, which is very close to the approximate value of $\frac{2}{3}\beta_{SPC}$. This translates into a system time about 70% greater than that of the TSP policy. The SFC policy, however, is much more efficient computationally.

Note that characterizing the performance of a given policy in this case reduces to estimating a single constant. Thus, by doing careful simulations over a restricted range of parameters to estimate γ, we obtain a closed form approximation for the behavior of a policy that can in turn be applied over a much wider range of operating conditions. This is to be contrasted with the usual case in simulations were a policy's behavior is usually characterized only with respect to a specific set of examples and little generalization beyond this set of examples is possible.

3.3.6 The Nearest Neighbor Policy

The last policy we consider is to serve the closest available demand after every service completion (nearest neighbor (NN) policy). The motivations for considering such a policy are: (1) the nearest neighbor was used in the heavy traffic lower bound of Theorem 3.2, and (2) the shortest processing time (SPT) rule is known to be optimal for the classical $M/G/1$ queue [15]. 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 only suggestive.

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

$$E[d_i|N_T] \leq \gamma_{NN}\frac{\sqrt{A}}{\sqrt{N_T}}, \tag{3.29}$$

where N_T is the number of demands in the system at a completion epoch, then by using a modification of the argument in [29] Section 5.5, it is possible to show that

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

where T_{NN} denotes the system time of the NN policy. The assumption (3.29) is analogous to Lemma 3.1 but unlike Lemma 3.1 has not been established formally.

We therefore performed simulation experiments identical to those for the SFC policy to verify the asymptotic behavior of T_{NN} and estimate γ_{NN}. The results showed that γ_{NN} is approximately 0.64. (See Figure 3.3.) This means that T_{NN} is about 10% lower than T_{SFC} but about 60% higher than the system time of the Modified TSP policy.

Policy (μ)		γ_μ	$\frac{T_\mu}{T^*}$
Light Traffic			
	FCFS	*	1.36
	SQM	*	1.00
Heavy Traffic			
	PART	$\sqrt{2c_1}$	7.35
	TSP	β	3.66
	SFC	0.66	3.08
	NN	0.64	2.89
	MOD TSP	$\frac{\beta}{\sqrt{2}}$	1.83

Table 3.1: Summary of Constant Values and guarantees for the Single Uncapacitated Vehicle DTRP

The results again confirmed that the system time T_{NN} follows the $\frac{\lambda A}{v^2(1-\rho)^2}$ growth predicted by the lower bound in Theorem 3.2. Figure 3.3 clearly shows this highly linear relationship.

3.4 Summary of Single Uncapacitated Vehicle Performance Bounds

We breifly pause now to reveiw our various performance guarantees for the single uncapacitated vehicle DTRP. These guarantees are summarized in Table 3.1. The numerical guarantees are based on the value $\gamma \approx 0.376$ from Chapter 5. Constant values for the SFC and NN policy are based on the simulation results of the previous section.

3.5 A Numerical Example

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

Figure 3.4 shows system times as a function of ρ for the light traffic case. The lower bound is also included. Note that although the SQM policy is asymptotically optimal as $\rho \to 0$, it is quickly surpassed by the FCFS policy as ρ increases. This is due to the extra travel distance of the SQM policy, which hinders the policy as queueing sets in. Also note that both policies reach their saturation points for relatively low values of ρ. The SFC and NN policies were comparable to the FCFS policy in very light traffic, which is to be expected since they essentially behave like the FCFS policy in this case. For $\rho > 0.05$ the SFC and NN policies quickly surpass the FCFC and SQM policies. Notice that the NN policy consistently performed better than the SFC policy even in the light traffic cases.

The heavy traffic results are shown in Figure 3.5. Note that the curves have nearly identical shapes as one would expect from the $\frac{\lambda A}{v^2(1-\rho)^2}$ asymptotic behavior of each policy. (Only the constant of proportionality differs.) The graphs show the sharp increase in system time as the traffic intensity increases. The Modified TSP policy is the best in this case with the SFC and NN policies second, having about 60-70% greater system times. The ordinary TSP policy and especially the PART policy are less effective.

This example suggests that the modified TSP policy is quite effective for heavy traffic. The SFC and NN policies are fairly effective over a wide range of traffic intensities. Indeed, if one locates a depot at the median of the region $\mathcal{A}$ and treats it as a permanent "demand", then both the SFC and NN policies can be made to behave like the SQM policy as $\rho \to 1$. These policies have another distinct advantage. Namely, they are nonparametric; that is, system parameters are not

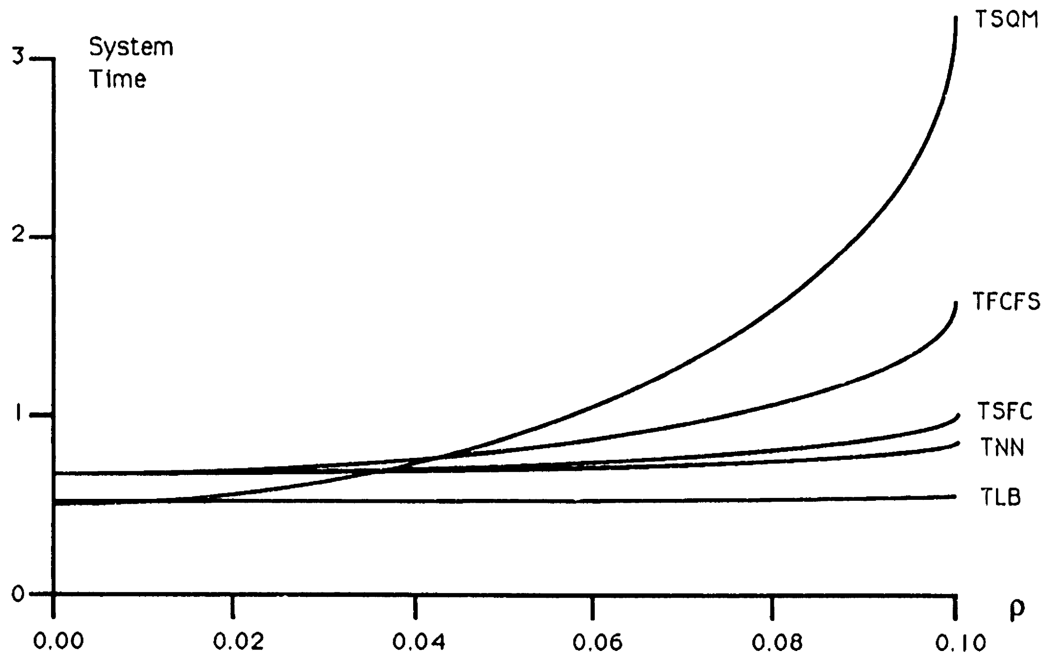

Figure 3.4: System Times for Light Traffic Case: Numerical Example

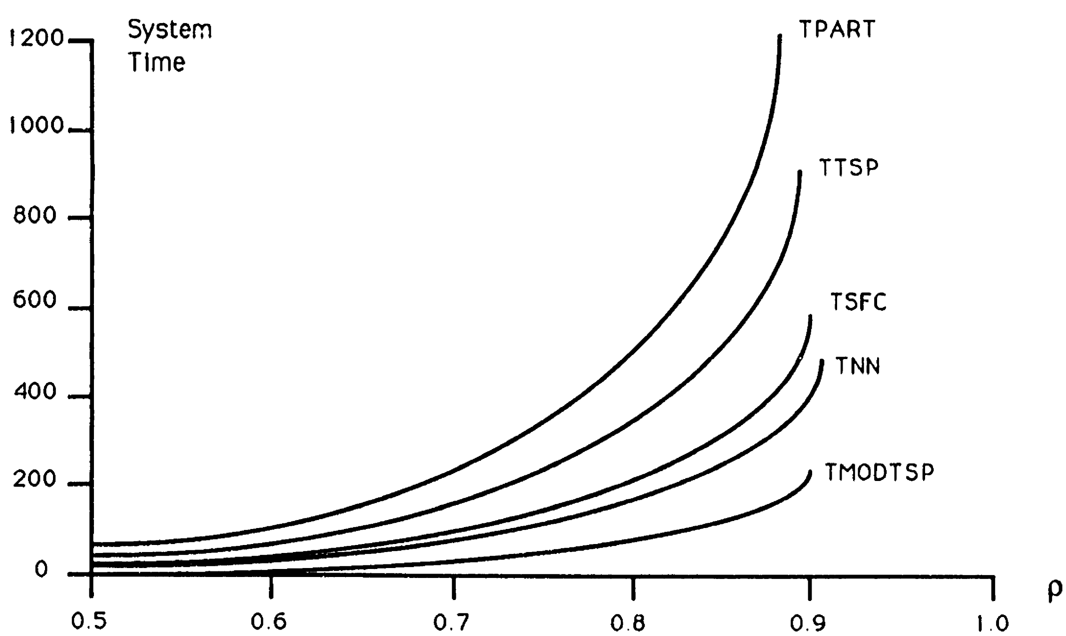

Figure 3.5: System Time for Heavy Traffic Policies: Numerical Example

needed to implement them as is the case for the TSP policies. This self regulating feature is especially desirable for system that operate under highly variable and/or unpredictable traffic conditions. Both policies are also both very computationally efficient. The combination of self regulating behavior and computational efficiency may make the SFC and NN policies attractive in practice despite the fact that their performance guarantees are not the best possible.

Chapter 4

The Multiple Capacitated Vehicle DTRP

In this chapter, we extend the basic results of Chapter 3 ultimately to the case where there are $m > 1$ identical vehicles with capacity constraints in the form of an upper bound of q on the total number of demands that can be serviced in any given trip from the depot. We begin in Section 4.1 by analyzing the m-vehicle case with $q = \infty$ through some relatively straightforward extensions of our earlier results. We then investigate the finite capacity case in Section 4.2.

4.1 The m-Vehicle, ∞-Capacity DTRP

4.1.1 Lower Bounds

A Light Traffic Lower Bound

The first bound is most useful in the case of light traffic ($\lambda \to 0$):

Theorem 4.1

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

Proof

We proceed as in the single-vehicle case and divide 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

$$T = W^d + W^s + \overline{s}, \tag{4.1}$$

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

As before, we bound W^d by noting 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, if we had the option of locating the m servers in the best *a priori* set of location, $\mathcal{D}^*$ this would certainly yield a lower bound on the expected distance between the nearest server and the demand's location. Hence,

$$W^d \geq \frac{1}{v} \min_{|\mathcal{D}|=m} E[\min_{x_0 \in \mathcal{D}} \|X - x_0\|]. \tag{4.2}$$

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 (4.1) and (4.2) establishes the theorem.

□ (Theorem 4.1)

Note that this is weaker than Theorem 3.1 for the case $m = 1$.

A Heavy Traffic Lower Bound

A lower bound useful for $\rho \to 1$ is provided by the following theorem, which generalizes the heavy traffic bound in Theorem 3.2:

Theorem 4.2 *There exists a constant γ such that*

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

where $\gamma \geq \frac{2}{3\sqrt{2\pi}} \approx 0.266$.

<u>Proof</u>

Theorem 4.2 is easily derived by modifying the proof of Lemma 3.1 slightly. Recall that $N_0 = |\mathcal{S}_0|$ is the number of demands in queue plus the number of servers in the system at the arrival epoch of our tagged demand. Thus, in the m-vehicle case, $E[N_0] = N + m$, and the bound on $E[Z^*]$ becomes

$$E[Z^*] \geq \gamma \frac{\sqrt{A}}{\sqrt{N + m/2}}. \tag{4.3}$$

Recall $\overline{d} \geq E[Z^*]$. By substituting these bounds into the stability condition,

$$\overline{s} + \frac{\overline{d}}{v} \leq \frac{m}{\lambda}, \tag{4.4}$$

we obtain

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

After rearranging, noting that $T = W + \overline{s}$ and $N = \lambda W$, we obtain the bound of Theorem 4.2.

□ (Theorem 4.2)

4.1.2 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.

Proposition 4.1

$$\frac{T_{mSQM}}{T^*} \longrightarrow 1 \qquad as \;\; \lambda \longrightarrow 0.$$

Proof

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 $\lambda \longrightarrow 0$, $\lambda_j \longrightarrow 0$ 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,

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

Comparing this to Theorem 3.1 establishes the proposition.

□ (Proposition 4.1)

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 4.2, it is clear that the mSQM policy has an unbounded cost relative to optimum for $\rho \longrightarrow \rho_c$ and certainly for $\rho \longrightarrow 1$.

4.1.3 Heavy Traffic Policies

We next turn our attention to heavy traffic ($\rho \rightarrow 1$) policies. We prove that policies based on randomized assignment of arrivals to servers have a constant factor guarantee for $\rho \rightarrow 1$, but this factor increases with m. We show for a version of the Modified TSP policy introduced in Chapter 3 that this dependence on m can be eliminated. Finally, we show that the same guarantees as in the single-vehicle case can be achieved if the service region $\mathcal{A}$ is divided into equal sized subregions and a single server heavy traffic policy is applied in each region.

Randomized Assignment (RA)

One possible strategy for a multiple-vehicle system is to allocate demands to vehicles using randomization. This policy, which we call randomized assignment ($RA\mu$), is defined as follows:

The $RA\mu$ Policy

Divide the Poisson input process into m Poisson sub-processes, one for each vehicle, using randomization. Assign one vehicle to service each sub-process using a heavy traffic, single-server policy μ.

Proposition 4.2

$$\frac{T_{RA\mu}}{T^*} \leq m\frac{\gamma_\mu^2}{\gamma^2} \qquad as \;\; \rho \rightarrow 1.$$

<u>Proof</u>

In this policy, each vehicle sees a demand arrival process with rate λ/m and operates independently in the entire region $\mathcal{A}$ to service it. The system time for randomized assignment is therefore simply

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

where as before $\rho \equiv \lambda\bar{s}/m$. Comparing this to the bound in Theorem 4.2 establishes the proposition.

□ (Proposition 4.2)

Observe that the performance guarantee for $RA\mu$ has the undesirable characteristic of increasing with the number of servers. Nevertheless, this shows that a sufficient condition for the existence of a stable policy is simply $\rho < 1$. This combined with the necessary condition $\rho < 1$ from Theorem 4.2 gives us the following theorem:

Theorem 4.3 *In the multiple, uncapacitated vehicle DTRP, a stable optimal policy exists if and only if* $\rho < 1$.

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. The policy is based on collecting customers into sets that can then be served using optimal TSP tours:

The G/G/m Policy

Let $\mathcal{N}_k$ denote the kth set of n demands to arrive, where n is a given constant that parameterizes the policy (e.g. $\mathcal{N}_1$ is the set of demands $1, \ldots, n$, $\mathcal{N}_2$ is the set of demands $n+1, \ldots, 2n$, etc.) Assume the server operates out of a depot at a random location in $\mathcal{A}$. As sets form, deposit them in a queue. Serve sets from the queue in FCFS order with the first available vehicle by following an optimal tour on their locations starting and ending at the depot. (The vehicle randomly selects one of the two possible orientations of these tours.) Optimize over n.

Proposition 4.3

$$\frac{T_{GGm}}{T^*} \le \frac{m+1}{2}\frac{\beta^2}{\gamma^2} \qquad as \;\; \rho \to 1.$$

Proof

Note that if one considers sets as customers, this policy defines a G/G/m queue. The interarrival distribution is Erlang 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, a random variable which we denote generically by τ, is the sum of the travel time around the tour, denoted L_n, and the n on-site service times. Thus, $E[\tau] = E[L_n]/v + n\bar{s}$ and $Var[\tau] = Var[L_n]/v^2 + n\sigma_s^2$, where $\sigma_s^2 = \overline{s^2} - \bar{s}^2$ is the variance of the on-site service time.

We next make use of a heavy traffic limit given in Equation (2.5). Letting W_{set} denote the waiting time of a set, this limit in our case gives

$$W_{set} \sim \frac{\frac{\lambda}{n}(\frac{n}{\lambda^2} + \frac{1}{m^2}(Var[L_n]/v^2 + n\sigma_s^2))}{2(1 - \frac{\lambda}{mn}(E[L_n]/v + n\bar{s}))} \tag{4.5}$$

$$= \frac{\lambda(1/\lambda^2 + \frac{1}{m^2}(\frac{var[L_n]}{nv} + \sigma_s^2))}{2(1 - \rho - \frac{\lambda}{m}\frac{E[L_n]}{nv})}. \tag{4.6}$$

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 asymptotic TSP results (*c.f.* Equations (2.10) and (2.10)) to assert that

$$\frac{E[L_n]}{n} \sim \beta\frac{\sqrt{A}}{\sqrt{n}}, \tag{4.7}$$

and

$$\frac{Var[L_n]}{n} \sim 0, \tag{4.8}$$

as $n \to \infty$. (Again, $\beta \approx 0.72$ is the Euclidean TSP constant.) Substituting these expressions above we obtain,

$$W_{set} \sim \frac{\lambda(1/\lambda^2 + \sigma_s^2/m^2)}{2(1 - \rho - \frac{\lambda}{m}\beta\frac{\sqrt{A}}{v\sqrt{n}})}. \tag{4.9}$$

For the queue to be stable, $\rho + \frac{\lambda}{m}\beta\frac{\sqrt{A}}{v\sqrt{n}} < 1$, which implies

$$n > \frac{\lambda^2\beta^2 A}{m^2v^2(1-\rho)^2}. \tag{4.10}$$

Therefore for $\rho \rightarrow 1$, n must indeed be large, and thus using asymptotic TSP results is justified. Also, as $\rho \rightarrow 1$, $\rho + \frac{\lambda}{m}\beta_{TSP}\frac{\sqrt{A}}{v\sqrt{n}} \rightarrow 1$ for all n satisfying (4.10), and thus we confirm the queue operates in heavy traffic.

As before, to get the system time of a demand we must add to (4.9) the time a customer waits for its set to form, W^-, and also the time it takes to complete service of the customer once the customer's set enters service, W^+. By conditioning on the position that a given customer takes within its set, one can show that

$$W^- = \frac{n-1}{2\lambda} \leq \frac{n}{2\lambda}.$$

Similarly,

$$W^+ \leq \frac{E[L_n]}{v} + \frac{1}{n}\sum_{k=1}^{n} k\overline{s} \leq \frac{\bar{\beta}\sqrt{nA}}{v} + \frac{n}{2}\overline{s},$$

where we have used the fact that the optimal tour on n points in a region of area A is bounded above by $\bar{\beta}\sqrt{nA}$ for some constant $\bar{\beta}$ (*c.f.* Equation (2.12)). Therefore, denoting the total system time by T_{GGm}, for $\rho \rightarrow 1$,

$$T_{GGm} \leq \frac{\lambda(1/\lambda^2 + \sigma_s^2/m^2)}{2(1-\rho-\frac{\lambda}{m}\beta\frac{\sqrt{A}}{v\sqrt{n}})} + \frac{n(1+m\rho)}{2\lambda} + \bar{\beta}\frac{\sqrt{nA}}{v}. \tag{4.11}$$

We would like to minimize (4.11) with respect to n to get the least upper bound. First, however, consider a change of variable to

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

Physically, y represents the ratio of the average travel time, $\overline{d} = \frac{\beta\sqrt{A}}{v\sqrt{n}}$ to its critical value $\frac{m(1-\rho)}{\lambda}$ (see equation (4.4)). With this change,

$$T_{GGm} \leq \frac{\lambda(1/\lambda^2 + \sigma_s^2/m^2)}{2(1-\rho)(1-y)} + \frac{\lambda\beta^2 A(1+m\rho)}{2m^2v^2(1-\rho)^2y^2} + \frac{\lambda\beta\bar{\beta}A}{mv^2(1-\rho)y}. \tag{4.12}$$

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

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

Substituting this approximation into (4.12) and noting that for $\rho \rightarrow 1$ the approximate y^* approaches 1 we have that as $\rho \rightarrow 1$

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

The leading term is proportional to $\frac{\lambda A}{m^2v^2(1-\rho)^2}$. Comparing this term to Theorem 4.2 establishes the proposition.

□ (Proposition 4.3)

We point out that the contribution due to the queueing term (W_{set}) is only $O((1-\rho)^{-3/2})$ and that the leading order term is due to W^- and the on-site service time component of W^+. Also, note that the leading term is still dependent on m but it increases like $(m+1)/2$ rather than m as in the randomized assignment case, which is clearly better but still somewhat unsatisfactory.

The Modified G/G/m Policy

A modification to the G/G/m policy can eliminate this dependence. The analysis requires Theorem 2.1 of Inglehart and Whitt [22] (see Chapter 2) on the behavior of the queue $\sum$GI/G/m: The modified G/G/m policy itself is defined as follows:

The Modified G/G/m Policy

For some fixed integer $k \geq 1$, divide $\mathcal{A}$ into k subregions of equal area using radial cuts centered at the depot (*i.e.* form k wedges of area A/k) . Within each region, form sets of size n/k as in the G/G/m policy and, as sets are formed, deposit them in a queue. Service the queue FCFS with the first available vehicle by following optimal tours as before. Optimize over n.

Proposition 4.4

$$\frac{T_{MOD\ GGm}}{T^*} \leq \frac{\beta^2}{2\gamma^2} \qquad as \ \ \rho \rightarrow 1.$$

Proof

Again, we will only sketch the proof of this proposition since the detailed analysis closely parallels that of the G/G/m policy. Observe the modified policy works very much like the G/G/m policy except that smaller tours are formed independently on subregions of $\mathcal{A}$.

This modification has the following effect on the three components of the system time, W^-, W^+ and W_{set} defined above: For W^-, which is the wait for a set to form, there is no change. This is because, although the number of demands in a set is reduced by $1/k$, the time between arrival of demands in a subregion is increased by a factor of k; hence, W^- remains the same. W^+, which is approximately one half of the service time of a set, is reduced by $1/k$ because the number of on-site services per tour is reduced by $1/k$ and, since both the area and number of points are reduced by $1/k$, the tour length $L_n \sim \beta\sqrt{nA}$ is also reduced by $1/k$.

This leaves the waiting time in queue for a set, W_{set}. The resulting queue is $\sum$GI/G/m since the input process in now the superposition of k independent renewal processes, one from each of the k subregions. Using Theorem 2.1, one can show that W_{set} again satisfies (4.9). Thus, combining the three terms W^-, W^+ and W_{set} and repeating the analysis, we obtain

$$T_{MOD\ G/G/m} \sim \frac{\beta^2}{2} \frac{\lambda A(1 + m/k)}{m^2 v^2 (1-\rho)^2}.$$

The proposition then follows by taking k to be arbitrarily large.

A D/G/m Version of the TSP Policy

We next briefly mention a D/G/m version of the TSP policy that has the same constant as the modified G/G/m policy. The D/G/m policy is again based on collecting demands into sets that can then be served using optimal TSP tours; however, sets are formed by clustering demands periodically in time and space as follows:

The D/G/m Policy

For some fixed integer $k \geq 1$, divide $\mathcal{A}$ into k subregions of equal area using radial cuts centered at the depot (*i.e.* wedges of area A/k). Number the regions $1, \ldots, k$ consecutively starting at an arbitrary region. At each time $t = (\frac{\theta}{k})j, \quad j = 1, 2, \ldots,$ form a set in subregion $(j \bmod k) + 1$ of all the demands that arrived in that subregion during $(t - \theta, t]$ (*i.e.* since the time the last set was formed in this subregion).

As sets are formed by this process, deposit them into a queue. Serve sets from the queue FCFS with the first available vehicle by following an optimal tour on the demands in the set starting and ending at the depot. (The vehicle randomly selects one of the two possible orientations of these tours.) Optimize over θ.

The behavior of this policy is summarized in the following proposition (The proof is omitted due to its similarity to the previous cases.):

Proposition 4.5

$$\frac{T_{D/G/m}}{T^*} \leq \frac{\beta^2}{2\gamma^2} \qquad as \ \rho \rightarrow 1.$$

To visualize the process, consider the case where $\mathcal{A}$ is a circle. The arm of a clock sweeps the circle $\mathcal{A}$ every θ time units, and, upon passing a subregion, deposits all the demands in that subregion into a set. The resulting sets are then served FCFS from a queue as in the G/G/m policy. In this way, sets are formed regularly every θ/k units of time and the number of demands in a set, N_θ, is a Poisson random variable with mean $(\lambda/k)\theta$.

This policy defines a D/G/m queue. The constant time between arrivals is θ/k, and the service time of sets are i.i.d. random variables. To analyze this queue, we again use the heavy traffic limit (2.5) and proceed as in the G/G/m case. The analysis of the policy, however, does not require Theorem 2.1 since the arrival process is deterministic.

Independent Partitioning Policies

The last policy we examine for the uncapacitated case is based on partitioning the service region:

The P_μ Policy

Divide the region $\mathcal{A}$ into m subregions of equal size. Assign one vehicle to each region, and have vehicles follow a single server policy μ to service demands that fall within their subregion.

Proposition 4.6

$$\frac{T_{P_\mu}}{T^*} \leq \frac{\gamma_\mu^2}{\gamma^2} \qquad as \quad \rho \rightarrow 1.$$

Proof

The effect of this independent partition is to reduce *both* the area and the arrival rate by a factor of $1/m$. Thus, it is immediate that

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

Comparing to the lower bound in Theorem 4.2 proves the proposition.

□ (Proposition 4.6)

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 interesting to compare the policy $P_{MOD\ G/G/1}$ to the modified $G/G/m$ policy. Assuming the same number of wedges, k, is used for both, the constant for $P_{MOD\ G/G/1}$ is $\beta\sqrt{\frac{1+1/k}{2}}$ while the constant for modified $G/G/m$ is $\beta\sqrt{\frac{1+m/k}{2}}$. While it is true that one can theoretically let k be an arbitrarily large constant, in practice the partitioning policy is to be preferred since for finite k it is always smaller.

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 4.2 is tight for all m, such a conclusion is premature; however, the idea seems highly plausible and is worth a conjecture:

Conjecture 4.1 *Let μ^* denote an optimal single-server DTRP policy in heavy traffic. Then P_{μ^*} is an optimal m-server policy in heavy traffic.*

4.2 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 can 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 depot, 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. We shall further assume $q > 1$, since otherwise the system behaves as an ordinary M/G/m queue.

4.2.1 A Heavy Traffic Lower Bound

We begin with the following lower bound:

Theorem 4.4 *For* $q > 1$,

$$T^* \geq \frac{\gamma^2}{9} \frac{\lambda A(1+\frac{1}{q})^2}{m^2 v^2 (1-\rho-\frac{2\lambda \bar{r}}{mvq})^2} - \frac{\bar{s}(1-2\rho)}{2\rho}.$$

where $\gamma \geq \frac{2}{3\sqrt{2\pi}} \approx 0.266$.

Proof

Consider demand i and the tour of length c_i that contains it. Randomly and independently select two *distinct* points in this tour and denote them j_1 and j_2. (Note that $j_1 = j_2$ is not allowed.) 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,

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

where Z_j^* is the distance to the nearest neighbor of j (i.e. Z^* as defined in the proof of Lemma 3.1). Since the points j_1 and j_2 are equally likely to be any of 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 Z_j^* above is the same as Z_i^*. Therefore, taking expectations on both sides, letting $i \rightarrow \infty$ and noting that $j_1 \rightarrow \infty$ and $j_2 \rightarrow \infty$ as $i \rightarrow \infty$ we obtain

$$\begin{aligned} \bar{c} &\geq 2\bar{r} + E[\sum_{j=j_*+1}^{j^*} Z_j^*] \\ &= 2\bar{r} + E[|\Delta j|]E[Z^*], \end{aligned} \tag{4.13}$$

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 Z_j^*. We next need the following lemma:

Lemma 4.1 $E[|\Delta j|] = \frac{q}{3}(1 + \frac{1}{q})$.

<u>Proof</u>

First, consider selecting points j_1 and j_2 that are *not* distinct (i.e $j_1 = j_2$ is allowed). The random variable Δj in this case is distributed as the difference between two independent, equiprobable selections from the set $\{1, 2, \ldots, q\}$. By considering the joint sample space, one can 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 implies that $E[|\Delta j|] = \frac{1}{3}(q - \frac{1}{q})$. Now, if we discard outcomes with $j_1 = j_2$, which occur with probability $\frac{1}{q}$, the probabilities of the remaining outcomes are scaled up by a factor of $1/(1 - \frac{1}{q})$. Since $j_1 = j_2$ outcomes contribute nothing to $E[|\Delta j|]$ above, it therefore follows that when selecting distinct points

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

□ (Lemma 4.1)

Using Lemma 4.1 and noting that lower bound on $E[Z^*]$ from Equation (4.3) applies in the capacitated case as well, (4.13) becomes

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

Using the fact that $\overline{d} = \overline{c}/q$ implies

$$\overline{d} \geq \frac{2\overline{r}}{q} + (1 + \frac{1}{q})\frac{\gamma\sqrt{A}}{3\sqrt{N + m/2}}.$$

Substituting this into the stability equation (4.4), rearranging and noting that $N = \lambda W$ and $T = W + \overline{s}$ we obtain the bound in Theorem 4.4.

□ (Theorem 4.4)

A few comments on this bound are in order. Note that if we take $q \to \infty$, the constant value is one third of the value in the uncapacitated case. This is somewhat troubling since we know that for $q = \infty$ the two problems are in fact equivalent. It is therefore worth exploring, briefly, the relevance of this bound.

First, note that there will always be a sufficiently large value of q (specifically, $q > \frac{3\lambda\overline{r}}{1-\rho}$) for which the uncapacitated bound in Theorem 4.2 dominates the capacitated bound in Theorem 4.4. Thus, when one views q as the independent parameter, Theorem 4.4 can often be irrelevant.

It is quite relevant, however, if we consider the region geometry, on-site service statistics and vehicle capacity q as given and view the arrival rate, λ, as the independent parameter. In this case one is typically interested in how the system behaves as the traffic rate increases toward its maximum value. Theorem 4.4 shows that a necessary condition for stability is

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

Thus, λ increasing toward its maximum value is equivalent to $\rho + \frac{2\lambda\overline{r}}{mvq} \to 1$. As the traffic intensity approaches this limit, the capacitated bound always dominates the uncapacitated bound; it is this asymptotic behavior that is well captured by Theorem 4.4.

The 1/3 factor appears to be mainly a by-product of the randomization used in the proof; our selection of two random points in effect "cuts out" one third of the local tour on the q points, which, when added to the two radial terms, forms the bound. The difficulty in eliminating the 1/3 factor is that one must bound the *sum* of the radial terms and the local tour terms. This is to be contrasted with the analysis of of the static VRP where typically only one of these terms dominates the cost. Indeed, this difficulty is related to the static VRP when $q = \Theta(n)$, for which not much is known (*c.f.* [20]).

We conjecture that the 1/3 value can be eliminated. Indeed, we know by Theorem 4.2 that for $q = \infty$ the performance guarantee is the same as in the uncapacitated case, and we show below in Proposition 4.9 that for $q = 2$ it is also the same.

As a further motivation, one can heuristically argue that

$$\bar{d} \geq 2\bar{r}\frac{1}{q} + E[Z^*](1 - \frac{1}{q})$$

as follows: a fraction $\frac{1}{q}$ of the arriving demands are first on the tour, for which d_i is the sum of two radial distances, one from the depot to the last demand in the previous tour and one from the depot to these first demands. The mean of this sum tends to $2\bar{r}$ in heavy traffic. (Exactly *how* it tends to $\bar{r}$ is the critical technical difficulty). The remaining points have d_i at least equal to the distance to the nearest point, Z_i, which gives us the above expression. Using this heuristic bound on $\bar{d}$ implies

$$T^* \geq \gamma^2 \frac{\lambda A(1 - \frac{1}{q})^2}{m^2 v^2 (1 - \rho - \frac{2\lambda\bar{r}}{mvq})^2} - \frac{\bar{s}(1 - 2\rho)}{2\rho}, \tag{4.14}$$

which, as we shall show below, would imply the same constant factor bound as in the uncapacitated case for all q. Finally, note that (4.14) is exact if we restrict ourselves to optimizing over the class of policies in which the radial connections to the depot have mean $\bar{r}$.

4.2.2 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 3.1 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.

4.2.3 Single-Depot Heavy Traffic Policies

We next construct two policies for the m-vehicle, q-capacitated problem for the case where all vehicles operate out of a single depot and $\rho + \frac{2\lambda\bar{r}}{mvq} \to 1$ (heavy traffic).

The Region Partitioning (qRP) Policy

The first heavy traffic policy is based on region partitioning and is defined as follows:

The qRP Policy

Divide the region $\mathcal{A}$ into k equal sized subregions (except perhaps on the boundary) using a square grid system centered at the depot as shown in Figure 4.1. When q consecutive demands arrive in a single subregion consider it the arrival of a set. Service sets in FCFS order by the first available vehicle 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 k.

Proposition 4.7

$$\frac{T_{qRP}}{T^*} \leq \frac{9\bar{\beta}^2}{2\gamma^2(1+\frac{1}{q})^2} \qquad as\ \rho \to 1.$$

where $\bar{\beta}$ is the TSP constant for q uniformly distributed points in a square (i.e. $\bar{\beta} = \frac{E[L_q]}{\sqrt{qA}}$).

<u>Proof</u>

We proceed as before and determine the waiting time for a set, W_{set}. Let L_i be length of the local, TSP tour containing the i-th customer and $E[L] = \lim_{i\to\infty} E[L_i]$.

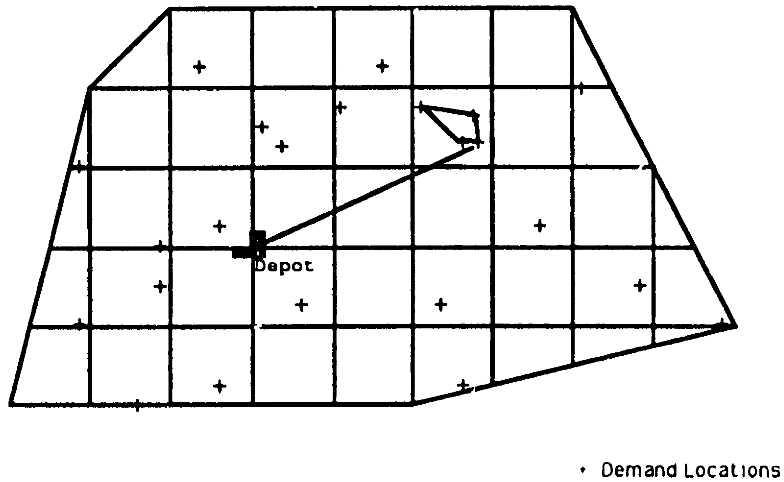

Figure 4.1: Subregions of the qRP Policy

Let $\bar{r}$ 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, $E[\tau]$, is given by

$$E[\tau] = 2\frac{\bar{r}}{v} + \frac{E[L]}{v} + q\bar{s}.$$

Denoting $Var[L]$ by σ_L^2, we have

$$Var[\tau] = 4\frac{\sigma_r^2}{v^2} + \frac{\sigma_L^2}{v^2} + q\sigma_s^2.$$

We point out that σ_s^2 is assumed finite, σ_r^2 is finite due to the boundedness of $\mathcal{A}$ and σ_L^2 is also finite (c.f. [31]).

Again, the queue formed by this policy is $\sum$GI/G/m. Thus by invoking Theorem 2.1 we obtain

$$W_{set} \sim \frac{\lambda(\frac{1}{\lambda^2} + \frac{1}{m^2}(\frac{4\sigma_r^2}{qv^2} + \frac{\sigma_L^2}{qv^2} + \sigma_s^2))}{2(1 - \frac{\lambda}{qm}(\frac{2\bar{r}}{v} + \frac{E[L]}{v} + q\bar{s}))}.$$

Define $\bar{\beta}$ to be the constant such that the length of the optimal tour on q uniform points in a square of a ea A satisfies $\frac{E[L]}{\sqrt{q}} = \bar{\beta}\sqrt{A}$. (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 $\bar{\beta}\sqrt{A}$ is an upper bound on $\frac{E[L]}{\sqrt{q}}$ (see [6]). Since each subregion has an area A/k, substituting this expression for $\frac{E[L]}{\sqrt{q}}$ above gives

$$W_{set} \sim \frac{\lambda(\frac{1}{\lambda^2} + \frac{1}{m^2}(\frac{4\sigma_{\bar{r}}^2}{qv^2} + \frac{\sigma_L^2}{qv^2} + \sigma_s^2))}{2(1 - \rho - \frac{2\lambda\bar{r}}{mqv} - \frac{\lambda\bar{\beta}\sqrt{A}}{mv\sqrt{qk}})},$$

where $\rho = \lambda\bar{s}/m$. Adding the expected wait for a set to form, which is at most $\frac{qk}{2\lambda}$, and the expected wait for service once the set enters service, which is at most $\frac{\bar{r}}{v} + \bar{\beta}\frac{\sqrt{qA}}{v\sqrt{k}} + \frac{q\bar{s}}{2}$, and making the change of variable

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

we obtain

$$T_{qRP} \leq \frac{\lambda\bar{\beta}^2 A}{2m^2v^2(1 - \rho - \frac{2\lambda\bar{r}}{mqv})^2 y^2} + \frac{\lambda(\frac{1}{\lambda^2} + \frac{1}{m^2}(\frac{4\sigma_{\bar{r}}^2}{qv^2} + \frac{\sigma_L^2}{qv^2} + \sigma_s^2))}{2(1 - \rho - \frac{2\lambda\bar{r}}{mqv})(1 - y)} + O(y).$$

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 $\rho + \frac{2\lambda\bar{r}}{mqv} \to 1$ by linearizing about the value $y = 1$. This yields

$$y^* \approx 1 - \frac{vm}{\bar{\beta}}\sqrt{\frac{(\frac{1}{\lambda^2} + \frac{1}{m^2}(\frac{4\sigma_{\bar{r}}^2}{qv^2} + \frac{\sigma_L^2}{qv^2} + \sigma_s^2))(1 - \rho - \frac{2\lambda\bar{r}}{mqv})}{2A}}.$$

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

$$T_{qRP} \sim \frac{\lambda\bar{\beta}^2 A}{2m^2v^2(1 - \rho - \frac{2\lambda\bar{r}}{mqv})^2}.$$

Comparing this leading term to the bound in Theorem 4.4 establishes the proposition.

□ (Proposition 4.7)

The qRP policy thus has a constant factor guarantee in the heavy traffic case. Note that this analysis has also established the sufficiency of the stability condition

$\rho + \frac{2\lambda\bar{r}}{mqv} < 1$ for the single-depot case since T_{qRP} is finite whenever this condition is satisfied. Thus, since Theorem 4.4 shows $\rho + \frac{2\lambda\bar{r}}{mqv} < 1$ is a necessary condition for stability, we have the following theorem:

Theorem 4.5 *In the multiple, capacitated vehicle DTRP with a single depot location, a stable optimal policy exists if and only if* $\rho + \frac{2\lambda\bar{r}}{mqv} < 1$.

The existence of such a policy also allows us to establish another rather intuitive theorem:

Theorem 4.6 *In the single-depot DTRP with vehicle capacities* $q > 1$, *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, $\bar{r}_{\mu^*} = \bar{r}^* + \Delta\bar{r}$ where $\Delta\bar{r} > 0$. Now consider the qRP policy with the depot located at the median. For notational convenience, define $\delta = 1 - \rho - \frac{2\lambda(\bar{r}^* + \Delta\bar{r})}{mvq}$ and $\epsilon = 2\lambda\Delta\bar{r}/mvq$. By our qRP results and Theorem 4.4, if μ^* is indeed optimal, then for all $\delta > 0$, T_{μ^*} must satisfy

$$\frac{\gamma^2\lambda A(1+\frac{1}{q})^2}{9m^2v^2\delta^2} - \frac{\bar{s}(1-2\rho)}{2\rho} \leq T_{\mu^*} \leq \frac{\bar{\beta}^2\lambda A}{m^2v^2(\delta+\epsilon)^2} + o(\delta+\epsilon)^{-3/2}.$$

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 4.6)

The Tour Partitioning (qTP) Policy

We next analyze a policy based on the tour partitioning (TP) scheme introduced by Haimovich and Rinnooy Kan [20] for the static vehicle routing problem. The policy is defined as follows:

The TP Policy

As in the G/G/m policy, collect demands into sets $\mathcal{N}_1, \mathcal{N}_2, \ldots$ of size n as they arrive and construct optimal tours on these sets. Starting at a randomly selected point in $\mathcal{N}_1$, split the tour into $l = \lceil n/q \rceil$ segments of q demands each (except, perhaps, for the last segment). Connect the end points of the segments to the depot to form l tours of at most q demands each. Assign the first available vehicle to service all the demands in the set using these tours. Repeat for $\mathcal{N}_2, \mathcal{N}_3, \ldots$ serving sets in FCFS order. Optimize over n.

Proposition 4.8

$$\frac{T_{qTP}}{T^*} \le \frac{9\beta^2(1+m)(1-\frac{1}{q})^2}{2\gamma^2(1+\frac{1}{q})^2} \qquad as \;\; \rho + \frac{2\lambda\bar{r}}{mqv} \to 1.$$

Proof

To analyze this policy, we need the mean and variance of the time to service a set. Since these set service times are i.i.d., it suffices to determine these quantities for the set $\mathcal{N}_1$. Let the random variable R_n denote the total radial connection distance for the set of tours on $\mathcal{N}_1$, L_n denote the length of an optimal tour on the set of points in $\mathcal{N}_1$ and $\tilde{L}_n$ denote the length of the portion of this tour that is actually used in the tour partition solution. The length of the total tour, denoted V_n, is therefore $V_n = R_n + \tilde{L}_n$.

To determine $E[V_n]$ we first condition on knowing the locations $\mathcal{N}_1 = \{X_1, \ldots, X_n\}$. As shown in [20], the sum of the lengths of the solutions produced by all of the n

possible starting points is

$$2l\sum_{i=1}^{n} r_i + (n-l)L_n.$$

Therefore it follows that the expected length of the tour obtained by randomly selecting one of these n starting points (still conditioned on $\mathcal{N}_1$) is simply

$$2l\frac{1}{n}\sum_{i=1}^{n} r_i + (1-\frac{l}{n})L_n.$$

Removing the conditioning on $\mathcal{N}_i$ we obtain

$$E[V_n] = 2l\overline{r} + (1-\frac{l}{n})E[L_n],$$

and therefore adding the on-site service times the expected time to service a set is

$$2l\frac{\overline{r}}{v} + (1-\frac{l}{n})\frac{E[L_n]}{v} + n\overline{s}. \tag{4.15}$$

The variance of the time to service a set is

$$\frac{Var[V_n]}{v^2} + n\sigma_s^2, \tag{4.16}$$

We shall evaluate $Var[V_n]$ shortly.

Using (4.15) and (4.16) in the G/G/m limit (2.5) and recalling that the mean and variance of the interarrival time of sets are n/λ and n/λ^2, we obtain the following limit for the waiting time for sets in queue, W_{set},

$$W_{set} \sim \frac{\lambda(\frac{1}{\lambda^2} + \frac{1}{m^2}(\frac{Var[V_n]}{nv^2} + \sigma_s^2)}{2\left(1-\rho-\frac{\lambda}{m}(\frac{2l\overline{r}}{vn} + (1-\frac{l}{n})\frac{E[L_n]}{vn}\right)}. \tag{4.17}$$

For $\rho + \frac{2\lambda\overline{r}}{mvq} \longrightarrow 1$ we show below that n must be large, therefore $E[L_n]/n \sim \beta\sqrt{A}/\sqrt{n}$ and $l/n \sim \frac{1}{q}$. For $Var[V_n]/n$, note that $Var[V_n] = Var[\tilde{L_n}] + Var[R_n] + 2Cov[\tilde{L_n}, R_n] \leq Var[\tilde{L_n}] + Var[R_n] + 2\sqrt{Var[\tilde{L_n}]Var[R_n]}$. It can be shown that $Var[\tilde{L_n}]/n \sim 0$ and $Var[R_n]/n \sim \sigma_r^2/q$; therefore, dividing through by n we have that $Var[V_n]/n \sim \sigma_r^2/q$. Substituting these above and adding the wait for a demand's set to form and the wait for a demand to be served once its set enters service we obtain

$$T_{qTP} \sim \frac{\lambda\left(\frac{1}{\lambda^2} + \frac{\sigma_s^2}{m^2} + \frac{\sigma_r^2}{m^2v^2q}\right)}{2\left(1-\rho-\frac{2\lambda\overline{r}}{mvq} - (1-\frac{1}{q})\frac{\lambda\beta\sqrt{A}}{mv\sqrt{n}}\right)} + \frac{n}{2\lambda}\left(1 + m(\rho + \frac{2\lambda\overline{r}}{mqv})\right) + \beta\sqrt{nA}.$$

Note that the stability condition is

$$\rho + \frac{2\lambda\overline{r}}{mvq} + (1 - \frac{1}{q})\frac{\lambda\beta\sqrt{A}}{mv\sqrt{n}} < 1,$$

which implies

$$n > \frac{(1-\frac{1}{q})^2\lambda^2\beta^2 A}{m^2v^2(1-\rho-\frac{2\lambda\overline{r}}{mvq})^2}.$$

Therefore assuming $n \to \infty$ as $\rho + \frac{2\lambda\overline{r}}{mvq} \to 1$ is consistent.

Making a change of variable to

$$y = \frac{(1-\frac{1}{q})\lambda\beta\sqrt{A}}{mv(1-\rho-\frac{2\lambda\overline{r}}{mvq})\sqrt{n}},$$

and using the approximation

$$y^* \approx 1 - \frac{mv}{\beta}\sqrt{\frac{(\frac{1}{\lambda^2} + \frac{\sigma_s^2}{m^2} + \frac{\sigma_r^2}{m^2v^2q})(1-\rho-\frac{2\lambda\overline{r}}{mvq})}{2(1-\frac{1}{q})A}},$$

we finally obtain that for $\rho + \frac{2\lambda\overline{r}}{mvq} \to 1$

$$T_{qTP} \sim \frac{\lambda\beta^2 A(1-\frac{1}{q})^2(1+m)}{2m^2v^2(1-\rho-\frac{2\lambda\overline{r}}{mvq})^2}. \tag{4.18}$$

Comparing this expression to the lower bound establishes the proposition.

□ (Proposition 4.8)

The presence of the factor $1+m$ can be eliminated using the following modified version of the qTP policy.

The Modified Tour Partitioning (MOD qTP) Policy

For some fixed integer $k \geq 1$, divide $\mathcal{A}$ into k subregions of equal area using radial cuts centered at the depot. Within each region, form sets of size n/k and form collections of feasible tours on these sets as in the qTP policy. As sets are formed, deposit them in a queue. Service the queue FCFS with the first available vehicle by following the collection of tours. Optimize over n.

The performance of this policy is described by the following proposition, which we shall not prove since the argument is only a slight modification of the previous proof:

Proposition 4.9

$$\frac{T_{MOD\ qTP}}{T^*} \leq \frac{9\beta^2(1-\frac{1}{q})^2}{2\gamma^2(1+\frac{1}{q})^2},$$

or equivalently,

$$T_{MOD\ qTP} \sim \frac{\lambda\beta^2 A(1-\frac{1}{q})^2}{2m^2v^2(1-\rho-\frac{2\lambda\overline{r}}{mvq})^2}.$$

This tour partitioning policy has several advantages over the qRP policy. First, note that it has a $(1-\frac{1}{q})^2$ factor multiplying its leading term. For low values of q, this improves the performance guarantee. Indeed, for $q=2$ we have

$$\frac{T_{MOD\ qTP}}{T^*} \leq \frac{\beta^2}{2\gamma^2},$$

which is the same as the best guarantee for the uncapacitated vehicle case. Comparing the leading behavior of $T_{MOD\ qTP}$ above to Equation (4.14), we see that the $\frac{\beta^2}{2\gamma^2}$ guarantee is also valid for all q if we restrict ourselves to optimizing over the class of policies which have a mean radial connection cost of $\overline{r}$.

Second, observe that the constant is the asymptotic value β *for all* values of q where as the qRP policy only achieves the asymptotic value β for large values of q. (Recall we used an upper bound $\overline{\beta}$ on β in the qRP policy.) This stems from the fact that, as the traffic intensity increases, the qRP policy reduces travel distance by forming optimal tours of q points on increasingly smaller subregions; the tour partitioning policies, by contrast, split an increasingly large tour on the entire region. Thus, the qRP policy constant is always based on finite tours of size q, while the tour partitioning policies achieves the asymptotic value β for any q. For these two reasons, we consider the tour partitioning policy to be superior to the region partitioning policy.

4.2.4 Heavy Traffic Policies for Some Symmetric Multi-Depot Cases

We now briefly describe some multi-depot cases 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. qRP with p vehicles) 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 [19] 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 $\bar{r}$, if follows that assigning p vehicles to each of the k medians is again provably good. Also, if one has the option of choosing k and p in this case, then $k = m$ and $p = 1$ are optimal since this choice minimizes $\bar{r}$, which in turn minimizes T.

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

Policy (μ)		γ_μ	$\frac{T_\mu}{T^*}$
Light Traffic			
	mSQM	*	1.00
Heavy Traffic ($q=\infty$)			
	G/G/m	$\sqrt{\frac{m+1}{2}}\beta$	1.83(m+1)
	RA_{MODTSP}	$\frac{\beta}{\sqrt{2}}$	1.83(m)
	MOD G/G/m	$\frac{\beta}{\sqrt{2}}$	1.83
	D/G/m	$\frac{\beta}{\sqrt{2}}$	1.83
	P_{MODTSP}	$\frac{\beta}{\sqrt{2}}$	1.83
Heavy Traffic ($q<\infty$)			
	qRP	$\frac{\bar{\beta}}{\sqrt{2}}$	8.03
	qTP	$\beta(1-\frac{1}{q})\sqrt{\frac{m+1}{2}}$	1.83(m+1)
	MOD qTP	$\beta\frac{1-\frac{1}{q}}{\sqrt{2}}$	1.83

Table 4.1: Summary of Constant Values and Guarantees for the Multiple Capacitated Vehicle DTRP

radial travel cost of $\bar{r}$ yet the arrivals must be evenly allocated to vehicles to achieve a uniform rate of λ/m. More sophisticated bounds and/or policies are probably needed in these cases.

4.3 Summary of Single Uncapacitated Vehicle Performance Bounds

The guarantees for the various policies from this chapter are summarized in Table 4.1. The numerical guarantees are again based on the value $\gamma \approx 0.376$ from Chapter 5. For the capacitated vehicle policies, the numerical values given are for the best case of $q=2$.

Chapter 5

The DTRP with General Distributions

In this chapter we extend our analysis to problems in which demand locations are distributed according to a continuous density $f(x)$ defined over $\mathcal{A}$ and arrivals are generated according to a renewal process with mean λ^{-1}, finite variance σ_a^2 and Laplace transform $A^*(s)$. The notation $f(x)$ is short for $f(x_1, x_2)$ ($x = [x_1\ x_2]$). Likewise, we write $\int f(x)dx$ for $\int\int f(x_1, x_2)dx_1 dx_2$. The function $f(x)$ satisfies

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

and

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

In our analysis, we will need to assume that this density satisfies $0 < \underline{f} \leq f(x) \leq \overline{f} < \infty, \forall\ x \in \mathcal{A}$. This assumption is needed for purely technical reasons, as we shall see shortly.

In this chapter we reverse conventions and define d_i to be the distance traveled from demand i to demand $i+1$; that is, the distance traveled after *departing* from i rather than the distance traveled to i. We shall also assume that decisions are taken only at service completion epochs and consist of choosing one of the demands in the system to visit next or perhaps choosing to visit a fixed depot location x_0. This assumption implies that only at x_0 can the vehicle choose the option of staying put. This property will turn out to be important in our analysis. We point out that in the uniform demand case, we allowed for a slightly more general class of policies in which vehicles can wait at any location x and change destinations at any point in time.

In turns out that in this general demand case we also need to distinguish between policies that provide the same level of service (*i.e.* same mean waiting time) for all locations, which we call *spatially fair* policies, and those which may produce waiting times that vary with location, which we call *spatially discriminating* policies. More formally:

Definition 5.1 *A policy μ is called spatially fair if*

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

and

Definition 5.2 *A policy μ is called spatially discriminatory if there exists sets $\mathcal{S}_1, \mathcal{S}_2 \subseteq \mathcal{A}$ such that*

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

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

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

As before, $T = W + \bar{s}$ for both fair and discriminatory policies.

5.1 Heavy Traffic Lower Bounds

In this section, we derive our main lower bounds. The bounds are established using a different proof technique than that used in Chapters 3 and 4 that not only allows us to consider general spatial distributions and arrival processes, but also improves on the constant value for the Poisson, uniform case ($f(x) = 1/A$).

5.1.1 Preliminary Definitions and Lemmas

Before beginning, some definitions are needed. For a given $\mathcal{S} \subseteq \Re^2$, let $N(\mathcal{S})$ denote the time average number of customers in queue located in $\mathcal{S}$. Note that since all demands are located in $\mathcal{A}$, $N(\mathcal{S}) = N(\mathcal{S} \cap \mathcal{A})$ for all sets $\mathcal{S}$. In particular, the time average number in queue $N = N(\mathcal{A})$ and if $\mathcal{S}$ is not a subset of $\mathcal{A}$ then $N(\mathcal{S}) = 0$

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

$$\phi(x) = \frac{1}{N} \lim_{z \to 0} \frac{N(\mathcal{C}(x,z))}{\pi z^2}. \tag{5.1}$$

We shall assume that if the system is stable, then this limit exists. Further, we will need that if $\underline{f} \leq f(x) \leq \overline{f}$ for all $x \in \mathcal{A}$, then $\underline{\phi} \leq \phi(x) \leq \overline{\phi}$ for all $x \in \mathcal{A}$ where these bounds on $\phi(x)$ may depend on $f(x)$ but not on N. Intuitively, this condition says that the density of demands in queue at any location x in the system grows essentially uniformly as $N \to \infty$ since $N\underline{\phi} \leq N\phi(x) \leq N\overline{\phi}$. We show below that this condition holds a fortiori. In particular, this assumption excludes policies that leave a customer permanently in the system at location x, since otherwise $N(\mathcal{C}(x,\delta)) \geq 1$ for all $\delta > 0$ and thus $\phi(x)$ is unbounded.

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

$$N(\mathcal{S}) = N \int_{\mathcal{S}} \phi(x)dx. \tag{5.2}$$

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

$$\int_{\mathcal{A}} \phi(x)dx = 1 \tag{5.3}$$

$$\phi(x) \geq 0 \quad \forall \ x \in \mathcal{A}. \tag{5.4}$$

We associate a queue with every subset $\mathcal{S}$, henceforth referred to as simply as the *the queue* $\mathcal{S}$, by considering $\mathcal{S}$ to be a "black box" that has arrivals (demands arriving to $\mathcal{S}$) and departures (service completions within $\mathcal{S}$). Let the random variable $Y(\mathcal{S})$

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

$$F^*_{Y(\mathcal{S})}(s) = \frac{A^*(s)p(\mathcal{S})}{1 - A^*(s)(1 - p(\mathcal{S}))}. \tag{5.5}$$

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

Lemma 5.1 *Let $||\mathcal{S}||$ denote the area of $\mathcal{S}$, then*

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

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

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

Proof

Note that if the region has Poisson arrivals, then for all $\mathcal{S}$, $E[n^+(\mathcal{S})] = N(\mathcal{S})$. This follows from PASTA [48] and the fact that customers are served sequentially (one at a time). Thus, to prove the lemma it is sufficient to show that the normalized interarrival time

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

has an exponential distribution for $||\mathcal{S}|| \to 0$, since this shows the arrival process is asymptotically Poisson and the above PASTA result applies. Letting $F^*_{\hat{Y}(\mathcal{S})}(s)$ denote the transform of $\hat{Y}(\mathcal{S})$ and suppressing the argument $\mathcal{S}$ in $p(\mathcal{S})$ for brevity we therefore have

$$F^*_{\hat{Y}(\mathcal{S})}(s) = \int_0^\infty e^{-st} dF_{\hat{Y}(\mathcal{S})}(t)$$

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

Note that $\|\mathcal{S}\| \longrightarrow 0$ implies $p(\mathcal{S}) \longrightarrow 0$. Therefore using (5.5), taking the limit as $p \longrightarrow 0$ and applying L'Hospital's rule we obtain

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

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

□ (Lemma 5.1)

The key insight shown by this lemma is that sampling a renewal process with low probability generates a Poisson process. Thus, for small regions $\mathcal{S}$, the arrival process to the queue $\mathcal{S}$ is approximately Poisson. We now apply this result to derive an important Lemma. Let Z_i^* denote the distance from the server to either the depot or the closest unserved demand (which ever is smaller) at the completion epoch of demand i. That is, Z_i^* is the decision that minimizes d_i; thus, $E[Z_i^*] \leq E[d_i]$ and

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

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

Lemma 5.2

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

<u>Proof</u>

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

$$P(Z^* \leq z | X = x) = P(n^+(\mathcal{C}(x,z)) > 0) \leq E[n^+(\mathcal{C}(x,z))] \tag{5.6}$$

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

Considering the service completion of our tagged demand as a departure from the queue $\mathcal{C}(x,z)$, we therefore have by Lemma 5.1 above that as $z \to 0$,

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

Expressing the integral above in terms of its asymptotic ($z \to 0$) value and substituting into the bound (5.6) implies

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

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

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

As mentioned, this bound is valid as long as the depot at x_0 is not within a radius $c^{-1/2} = [N\pi\phi(x)]^{-1/2}$ of the location x. Let $\mathcal{D}(N) = \{x | \, \|x - x_0\| \leq [N\pi\phi(x)]^{-1/2}\}$ denote the set of points for which the bound is not valid. We next establish that for large N, the contribution to the lower bound from the set $\mathcal{D}(N)$ is negligible

– and it here that we need our technical assumptions on $f(x)$ and $\phi(x)$. First note that $\phi(x) \geq \underline{\phi}$ $\forall N$ implies $\int_{\mathcal{D}(N)} dx \leq O(1/N)$. Using the trivial bound $P(Z^* > z|X = x) \geq 0$ for the points in $\mathcal{D}(N)$ and removing the conditioning $\{X = x\}$ implies,

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

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

□ (Lemma 5.2)

Lemma 5.2 can be used to prove the following intermediate result:

Lemma 5.3 *There exists a constant γ such that*

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

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

Proof

Consider the following necessary condition for stability

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

Using the fact that $E[Z^*] \leq \bar{d}$, multiplying the second term on the left hand side above by $\frac{\sqrt{N}}{\sqrt{N}}$ and rearranging implies

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

Note that since N is at least as large as the mean number in queue in the corresponding G/G/m queue (*i.e.* the queue with $v = \infty$), as $\rho \rightarrow 1$, we must have $N \rightarrow \infty$. Therefore taking the limit as $\rho \rightarrow 1$ (and consequently $N \rightarrow \infty$) on both sides above and applying Lemma 5.2 we obtain

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

where $\gamma \geq \frac{2}{3\sqrt{\pi}}$. Squaring both sides and using $T \geq W = \frac{N}{\lambda}$ we obtain Lemma 5.3.

□ (Lemma 5.3)

5.1.2 A Spatially Fair Lower Bound

As mentioned, Lemma 5.3 is only an intermediate result since the functions $\phi(x)$ remains unspecified. Determining $\phi(x)$ will give us our main heavy traffic theorems.

Theorem 5.1 *Within the class of spatially fair policies*

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

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

Before proving this theorem, note that it differs from Theorem 4.1 in that it is an asymptotic bound while Theorem 4.1 is valid for all values of ρ; however, it improves on the previous bounds since the constant value is larger by a factor of $\sqrt{2}$.

<u>Proof</u>

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

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

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

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

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

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

Letting $z \to 0$ above and noting that the above equality is true for all sets $C(x, z)$ implies that $\phi(x) = f(x) \ \forall \ x \in \mathcal{A}$. Making the substitution $\phi(x) = f(x)$ in the bound in Lemma 5.3 we obtain the theorem.

□ (Theorem 5.1)

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

5.1.3 A Spatially Discriminatory Lower Bound

Theorem 5.1 gives an asymptotic bound for the case where fairness is a constraint, perhaps imposed as a matter of policy. What is the system time behavior when this constraint is relaxed? The answer, in part, is provided by our second main theorem:

Theorem 5.2 *Within the class of spatially discriminatory policies*

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

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

<u>Proof</u>

Since no assumption of fairness is made, consider the following minimization problem for the integral term in Lemma 5.3.

$$\begin{aligned} z^* &= \min \int_{\mathcal{A}} \phi^{-1/2}(x) f(x) dx \qquad (5.7) \\ \text{subject to} \quad & \int_{\mathcal{A}} \phi(x) dx = 1 \\ & \phi(x) \geq 0. \end{aligned}$$

Using the value z^* as a lower bound on the integral term in Lemma 5.3 will give us Theorem 5.1. Note that the objective function is convex in $\phi(x)$ and the constraints are linear; thus, (5.7) is a convex program.

Relaxing the equality constraint above with a multiplier, we obtain the following Lagrangian dual

$$\begin{aligned} z^*(\mu) &= \min_{\phi(x)\geq 0} \int_{\mathcal{A}} \phi^{-1/2}(x) f(x) dx + \mu \left[\int_{\mathcal{A}} \phi(x) dx - 1 \right] \\ &= \int_{\mathcal{A}} \min_{\phi(x)\geq 0} \left[\phi^{-1/2}(x) f(x) + \mu \phi(x) \right] dx - \mu \end{aligned} \tag{5.8}$$

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

$$-\frac{1}{2}[\phi^*(x)]^{-3/2} f(x) + \mu^* = 0 \quad \forall\, x \in \mathcal{A} \tag{5.9}$$

$$\int_{\mathcal{A}} \phi^*(x) dx = 1 \tag{5.10}$$

$$\phi^*(x) \geq 0 \tag{5.11}$$

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

$$\phi^*(x) = \left[\int_{\mathcal{A}} f^{2/3}(x) dx \right]^{-1} f^{2/3}(x) \tag{5.12}$$

$$\mu^* = \frac{1}{2} \left[\int_{\mathcal{A}} f^{2/3}(x) dx \right]^{3/2} \tag{5.13}$$

is such a pair. The fact that (5.7) is a convex program implies that these conditions are also sufficient to assure global optimality. Substituting the value $\phi^*(x)$ above into Lemma 5.3 gives us the theorem.

□ (Theorem 5.2)

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

5.2 Heavy Traffic Policies

We next examine two policies that have provably good performance with respect to the lower bounds of Theorems 5.1 and 5.2. The policies are modifications of policies introduced in Chapters 3 and 4.

5.2.1 A Provably Good Spatially Fair Policy

The spatially fair policy we consider is defined as follows:

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

The performance of this policy is given by the following proposition.

Proposition 5.1 *Let T_F^* be the optimal system time over the class of spatially fair policies. Then*

$$\frac{T_F}{T_F^*} \leq \frac{\beta^2}{2\gamma^2} \quad \textit{as } \rho \rightarrow 1.$$

Before proving this proposition, we note that since $\beta \approx 0.72$ and $\gamma \geq \frac{2}{3\sqrt{\pi}} \approx 0.376$,

$$\frac{T_F}{T_F^*} \leq 1.8.$$

Thus, this fair policy is guaranteed to within about 80% of the optimal policy in heavy traffic.

Proof

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

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

Observe that $kf(x)$ is the conditional density in any given subregion. From the asymptotic TSP results of Equation (2.11), we can therefore assert that almost surely

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

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

$$\begin{aligned}
\frac{E[\tau]}{(n/k)} &= \bar{s} + \frac{1}{v\sqrt{n}} \sum_{i=1}^{k} \frac{E[L_i]}{\sqrt{n}} \\
&\sim \bar{s} + \frac{1}{v\sqrt{n}} \sum_{i=1}^{k} \beta \int_{\mathcal{A}_i} f^{1/2}(x)dx \\
&= \bar{s} + \frac{\beta}{v\sqrt{n}} \int_{\mathcal{A}} f^{1/2}(x)dx \qquad (5.14)
\end{aligned}$$

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

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

Hence

$$\frac{\sigma_\tau^2}{(n/k)} = \sigma_s^2 + O(1). \qquad (5.15)$$

We will use (5.14) and (5.15) shortly.

Let $\hat{\lambda}_i = \lambda/n$ denote the arrival rate of sets to region i and $\hat{\lambda} = \sum_{i=1}^{k} \hat{\lambda}_i = \frac{k\lambda}{n}$ denote the overall arrival rate of sets. For a randomly chosen demand, let W^- denote the waiting time for its set to form, W_{set} denote its waiting time in queue and W^+ denote its wait to be serviced once its set enters service.

Note that as in the uniform demand case

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

and

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

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

We have by Theorem 2.1 that

$$W_{set} \sim \frac{\sum_{i=1}^{k} \hat{\lambda}_i^3 \sigma_{a_i}^2 + \frac{1}{m^2}(\frac{m}{E[\tau]})^3 \sigma_\tau^2}{2(\frac{m}{E[\tau]})^2(1 - \frac{\hat{\lambda}E[\tau]}{m})} \quad as \ \frac{\hat{\lambda}E[\tau]}{m} \longrightarrow 1.$$

Since the interarrival time in each subregion is a geometric sum of interarrival times in the entire region, one can easily show that the variance of the interarrival time of sets from subregion i, $\sigma_{a_i}^2$, is given by

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

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

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

We shall use the fact that for large values of k, the right hand side above is approximately $\frac{1}{\lambda}\hat{\lambda}^2$. In heavy traffic, $\hat{\lambda} \approx \frac{m}{E[\tau]}$. Using this fact and applying (5.14) and (5.15), we can therefore rewrite the above limit for W_{set} as

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

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

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

so $n \longrightarrow \infty$ as $\rho \longrightarrow 1$ and thus the above asymptotic analysis is valid in heavy traffic.

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

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

Making a change of variable to

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

the bound can be written

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

An approximate optimal value for y is

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

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

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

where the second order term is $O((1-\rho)^{-3/2}$. The proposition then follows comparing the above leading behavior to the bound in Theorem 5.1 and choosing k arbitrarily large.

□ (Proposition 5.1)

5.2.2 A Provably Good Spatially Discriminatory Policy for Piece-Wise Uniform Demand

We next propose a policy that achieves a performance guarantee of $\frac{\beta^2}{2\gamma^2}$ with respect to the spatially discriminatory bound when f is a *piecewise uniform* density, *i.e.* there exists a partition of $\mathcal{A}$ into J subsets $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_J$ such that $f(x) = \mu_j \ \forall\, x \in \mathcal{A}_j$, $j = 1, 2, \ldots, J$. In particular, for such a density note that

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

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

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

Let the the system time of this policy be denoted T_D. We shall prove the following proposition:

Proposition 5.2 *If f is a piecewise uniform density and T_D^* is the optimal system time over the class of discriminatory policies, then*

$$\frac{T_D}{T_D^*} \leq \frac{\beta^2}{2\gamma^2} \quad as\ \rho \rightarrow 1.$$

Before beginning the proof, we note that this is again the same guarantee 1.8 as in the fair case.

<u>Proof</u>

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

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

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

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

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

$$\begin{aligned}\frac{E[\tau]}{(n/k)} &\to \bar{s} + \frac{\beta}{v\sqrt{n}}\sum_{j=1}^{J} p_j\mu_j^{-1/3} \\ &= \bar{s} + \frac{\beta}{v\sqrt{n}}\sum_{j=1}^{J} \mu_j^{2/3} A_j \qquad (5.16)\end{aligned}$$

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

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

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

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

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

$$\frac{\sigma_\tau^2}{n/k} = \sigma_s^2 + \frac{O(1)}{k}. \qquad (5.17)$$

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

$$\begin{aligned}W^- &\le \sum_{j=1}^{J} p_j \frac{1}{2}\left(\frac{n/k}{(p_j\lambda)/k_j}\right) \\ &= \frac{n}{2\lambda}\sum_{j=1}^{J}\frac{k_j}{k} \\ &= \frac{n}{2\lambda}\sum_{j=1}^{J}\mu_j^{2/3}A_j \qquad (5.18)\end{aligned}$$

and

$$\begin{aligned} W^+ &\leq \frac{n}{2k}\bar{s} + \frac{1}{v}\sum_{j=1}^{J} p_j E[L_j] \\ &= \frac{n}{2k}\bar{s} + O(\sqrt{n}). \end{aligned} \tag{5.19}$$

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

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

Again, we use the fact that for large values of k, the right hand side above is approximately $\frac{1}{\lambda}\hat{\lambda}^2$. Substituting this approximate expression into Theorem 2.1 and using Equations (5.16) and (5.17) we obtain,

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

Adding the bound (5.18) and (5.19) to the above expression we obtain that as $\rho \to 1$,

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

In terms of

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

we have

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

An approximate optimal value for y is

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

Substituting this into the bound on T_D for $\rho \to 1$ we obtain

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

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

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

Comparing this to the lower bound in Theorem 5.2 establishes the proposition.

□ (Proposition 5.2)

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

5.3 A Numerical Investigation of the Performance of the SFC and NN Policies for Generally Demand Distributions

In this section, we briefly examine some simulation results for the space filling curve (SFC) and nearest neighbor (NN) policies as defined in Chapter 3 for the general demand distribution case. We show that the SFC policy acts approximately like a fair policy. The NN policy, on the other hand, appears to act neither like a pure fair

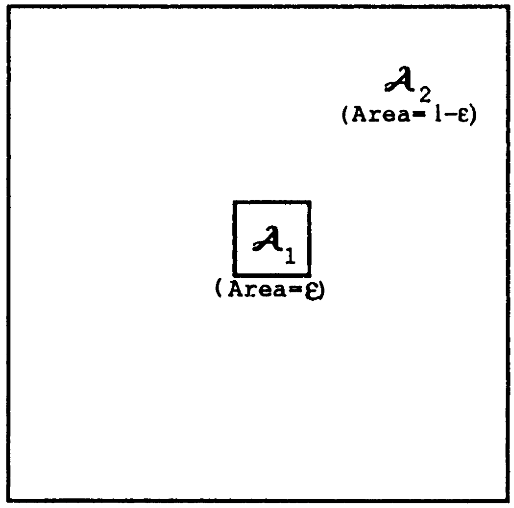

Figure 5.1: An Extreme Case General Demand Example

policy nor like an optimal discriminatory policies; rather, its performance seems to lie somewhere between these two extremes.

The general demand distribution used in the simulations is the one shown in Figure 5.1 The regions $\mathcal{A}_1$ and $\mathcal{A}_2$ have areas ϵ and $1-\epsilon$ respectively. Within each region demands are uniformly distributed. Pointc fall in region $\mathcal{A}_1$ with probability $1-\delta$ and in region $\mathcal{A}_2$ with probability δ. Thus, the density is piecewise uniform with

$$f(x) = \begin{cases} \frac{1-\delta}{\epsilon} & x \in \mathcal{A}_1 \\ \frac{\delta}{1-\epsilon} & x \in \mathcal{A}_2 \end{cases} \tag{5.20}$$

We used identical simulation techniques (*i.e.* same simulation code with different $f(x)$) as in Chapter 3. (See Chapter 3 for details.) To estimate the dependence of the system time for each policy, we set $\epsilon = 10^{-4}$ and fixed $A = 1$, $\bar{s} = 0.1$, $\sigma_s^2 = 0$ and $\rho = 0.8$. Then, a different simulation run was performed for eleven values of δ in the range 0.05 to 0.9999. (This last value corresponds to uniform demand.) The observed average number in the system (which is proportional to the average system time) was recorded for each δ for both the SFC and NN policies.

Before examining the results of these runs, it is useful to consider the following representation of the dependence of the system time on the density $f(x)$:

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

where

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

In the uniform case $\alpha = 0$, in the fair case $\alpha = 1/2$ and in the optimal discriminatory case $\alpha = 2/3$. For the particular density $f(x)$ given by (5.20) and for ϵ small,

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

and therefore for a particular policy μ

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

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

Figure 5.2 shows a log-log plot of the sample average number in the system as a function of δ for our simulation runs. The estimate of the slope of each line is shown in Figure 5.2 as well. For the SFC policy, the estimated slope of 0.80 corresponds to $\alpha = 0.44$ while for the SFC policy, the slope of 1.37 implies $\alpha = 0.58$. These values suggest that the SFC policy performs like a fair policy since its value of α is close to 1/2. (Though the performance appears to be somewhat worse than a purely fair policy.) The NN policy, on the other hand, seems to be somewhere between a fair and an optimal discriminatory policy; that is, it achieves a higher value of α than a fair policy could, but does not achieve as high a value of α as optimal discriminatory policies.

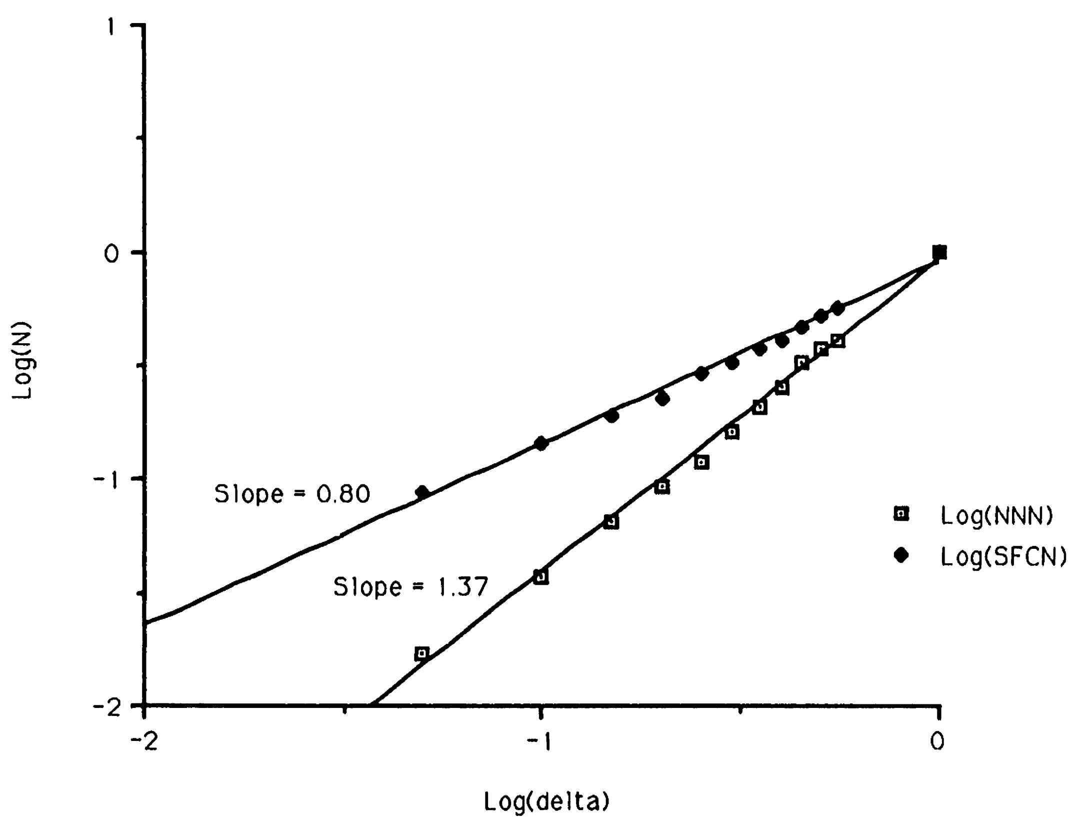

Figure 5.2: Simulation Results for SFC and NN Policies for General Demand Distribution

These results give the approximate behavior of the SFC and NN policies for general demand distributions. They also suggest a means of characterizing the behavior of other policies that cannot be analyzed mathematically; namely estimate γ_μ and α_μ as we did above and use the approximation

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

For example, this estimation might be performed using operating data from a "live" system, in which case the results could be used as a means of evaluating the effectiveness of a firms current operating practice.

5.4 Relationship Between Fair and Discriminatory Behavior

To review, we have determined that

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

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

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

Proposition 5.3 *For any continuous density function $f(x)$ defined over the region $\mathcal{A}$ of area A*

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

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

Proof

The proof requires the following inequality of Hardly, Littlewood and Pòlya [21]:

Lemma 5.4 (Hardly, Littlewood and Pòlya) *If $\alpha > 1$ or $\alpha < 0$, $g(x) \geq 0$ and $h(x) \geq 0$ then*

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

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

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

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

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

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

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

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

□ (Proposition 5.3)

Proposition 5.3 says that a uniform density is the worst possible and that any deviation from uniformity in the demand distribution will strictly lower the optimal mean system time in either the fair or discriminatory case. In addition, not requiring fairness in the service policy will result in a strict reduction of the optimal mean system time for any nonuniform distribution f. Also, note that when the density is uniform there is nothing to be gained by not providing fair service.

One may question how different the system times for a discriminatory and fair policy policy may be in general. That is, how much can one gain by discriminating according to location? Or, alternatively, how much does one lose by imposing a fairness constraint? The answer is that in the worst case the two can be arbitrarily far apart as illustrated by the simulation example in Figure of 5.1. For this density, it is straightforward to show that for a fixed $\delta > 0$ and $\epsilon \rightarrow 0$,

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

and

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

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

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

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

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

We can examine this phenomenon more formally. Recall that $N\phi(x)$ is the time average density of customers in queue and that $\phi(x)$ is proportional to $f^{2/3}(x)$ for optimal discriminatory policies. Let $W(x)$ be the waiting time at location x. By Little's Theorem, $W(x) = (\lambda f(x))^{-1}(N\phi(x))$, which implies that in an optimal discriminatory policy

$$\frac{W(x_1)}{W(x_2)} = \left[\frac{f(x_2)}{f(x_1)}\right]^{1/3}$$

for any two locations x_1 and x_2. In the example above with $x_1 \in \mathcal{A}_1$ and $x_2 \in \mathcal{A}_2$, this gives

$$\frac{W(x_1)}{W(x_2)} = \left[\frac{\delta\epsilon}{(1-\delta)(1-\epsilon)}\right]^{1/3},$$

which shows that for small values of ϵ and δ the discriminatory policy imposes a much greater waiting time on demands in $\mathcal{A}_2$ relative to those in $\mathcal{A}_1$. However, since only a fraction δ of the demands come from $\mathcal{A}_2$ the overall mean waiting time, $W = \delta W(x_1) + (1-\delta)W(x_2)$, is in fact reduced.

5.5 On the Tightness of the Lower Bounds for the General Case

In the proof of Lemma 5.2, one can see that very little of the vehicle routing "structure" inherent in the DTRP was used. Indeed, we only assumed that the service was sequential (*i.e.* one demand served at a time), which allowed us the establish that the mean number left behind by a departure from any given region was the same as the time average number in queue in that region. The bound therefore applies to any system in which points arrive randomly to a Euclidean region and are then removed sequentially according to some given rule. For example, we might remove a point after it spends a fixed amount of time τ in the system, in which case the expected nearest neighbor distance $E[Z^*]$ and the mean number in queue N would also satisfy Lemma 5.2. A DTRP policy, in this sense, simply defines one such rule for removing points; namely, remove a point after a vehicle following a given policy has completed its on-site service. In this section, we show that the lower bound in Lemma 5.2 is in fact tight within this broader class of "removal rules", and therefore more vehicle routing features of the DTRP need to be exploited to improve on our lower bounds.

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

Optimal Removal Rule

> Each arrival of a new point triggers a ***round*** of removals. A round of removals proceeds as follows: The oldest point in the system that is within a radius z of any neighboring point is removed. ($z > 0$ is an arbitrarily small constant.) The second oldest point with z of any of the remaining points is then removed, etc.. The round continues until no more points are left within z of any other point. Though these removals are sequenced, we assume the round of removals takes place instantaneously. This process is repeated for every arriving point.

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

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

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

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

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

Using the expression for W above we have

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

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

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

Comparing this to the bound in Lemma 5.2 shows that the lower bound is indeed tight within the class of sequential removal rules if points are uniformly distributed.

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

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

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

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

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

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

For the same reasons as in the uniform case,

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

Unconditioning implies

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

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

These results show that no improvement in the lower bound Lemma 5.2 is possible if we exploit only the fact that DTRP policies define removal rules. (That the bound is tight even within this broader class is still somewhat surprising given that the derivation of Lemma 5.2 used some seemingly crude arguments, such as bounding probabilities by expectations.) Thus, further improvements in lower bounding the value of $\bar{d}$ for the DTRP must necessarily use more of the vehicle routing characteristics of the problem. This bound is in essence a dynamic counterpart to the following static nearest neighbor bound for n uniformly distributed points in a region of area A:

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

which is used in the probabilistic analysis of such Euclidean problems such as the TSP, Matching and Minimum Spanning Tree [36]. In the same sense that this nearest neighbor bound is weak for the static TSP, one can see that the bound of Lemma 5.2 is likely to be weak for the DTRP. This suggests that the provable performance bound of 1.83 for the fair and discriminatory policies is too pessimistic. Indeed, we conjecture that these policies are asymptotically optimal.

5.6 General Demand Distributions and Capacitated Vehicles

Most of the results for the general demand distributions extend to the capacitated vehicle case as well. We shall only summarize results in this section since the analysis closely parallels the arguments we have seen in previous sections of this chapter and in Chapter 4.

By simply using the more general bound on the nearest neighbor distance $E[Z^*]$ of Lemma 5.2 in the arguments of Theorem 4.4, one can show the following theorems:

Theorem 5.3 *Within the class of spatially fair policies*

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

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

Theorem 5.4 *Within the class of spatially discriminatory policies*

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

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

A provably good fair policy for the finite capacity case can be obtained by modifying the fair policy from Proposition 5.1 as follows: as sets of size n/k are formed, partition these sets into feasible tours of at most q points using the tour partitioning heuristic of Haimovich and Rinnooy Kan [20] as was done for the Modified qTP policy of Proposition 4.9. Serve these sets FCFS and optimize over n. For large k, the resulting system time, T_{qF}, then satisfies

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

which implies the same performance guarantee as in the uniform case.

An identical tour partitioning modification applied to the sets formed in the spatially discriminatory policy of Proposition 5.2 gives a policy with a system time, T_{qD}, satisfying

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

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

Chapter 6

Further Extensions

6.1 Higher Dimensions

Most of the DTRP bounds and policies can be extended to Euclidean subsets $\mathcal{A}$ of $\Re^d$ for arbitrary dimension d. We examine this extension briefly in this section.

Consider first the uniform case. Let V denote the volume of $\mathcal{A} \in \Re^d$. Then repeating the proof of Theorem 4.2 for general d one can show the following theorem for the uncapacitated, m-server DTRP.

Theorem 6.1

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

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* $\Re^d$.

Similar modifications to the proofs of Theorems 5.1 and 5.2 give the following bounds:

Theorem 6.2 *Within the class of spatially fair policies*

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

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

Theorem 6.3 *Within the class of spatially discriminatory policies*

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

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

Note that $\gamma(d)$ is larger in these bounds by a factor of $2^{1/d}$.

Again, similar results 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

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

where $\Xi(d) = V$ for the uniform case, $\Xi(d) = \left[\int_{\mathcal{A}} f^{\frac{d-1}{d}}(x)dx\right]^d$ for the spatially fair case and $\Xi(d) = \left[\int_{\mathcal{A}} f^{\frac{d}{d+1}}(x)dx\right]^{d+1}$ for the discriminatory case.

For example, the modified TSP policy in d dimensions has a constant of $\gamma_{mTSP}(d) = \frac{\beta(d)}{2^{1/d}}$, where $\beta(d)$ is the d dimensional TSP constant.

An interesting result is found by examining this policy for $d \to \infty$. In [12], it was conjectured and subsequently proved in [41] that for $d \to \infty$

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

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

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

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

Theorem 6.4 *For the uncapacitated, m-server DTRP, the modified TSP policy is an optimal heavy traffic policy asymptotically as* $d \to \infty$.

This theorem gives further insight on the asymptotic optimality of the modified TSP policy, which we conjecture is also optimal for finite d as well.

6.2 Routing to Minimize Travel and Waiting Cost

We next turn our attention in a different direction and reconsider the objective function for our problem. Though we have concentrated throughout our discussion on minimizing system time, in many practical problems there is in fact a mixed objective involving waiting *and* travel costs. The value $\bar{d}$, the travel distance per demand served, is perhaps the most natural measure of the travel cost in our formulation since over an infinite time horizon the total travel distance is always infinite. Thus, rather than simply minimizing T we may in fact be interested in a more general objective function of the form

$$g(T, \bar{d}),$$

where g penalizes both the system time T and travel cost $\bar{d}$.

It turns out that this objective can be easily incorporated in the policies we have proposed. In our analysis, we consistently 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 $\bar{d}/v$ 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 restrict ourselves to the case of a single, uncapacitated vehicle (*i.e.* the TSP policy as defined in the Chapter 3) to illustrate the tradeoff. Similar results apply for the other cases.

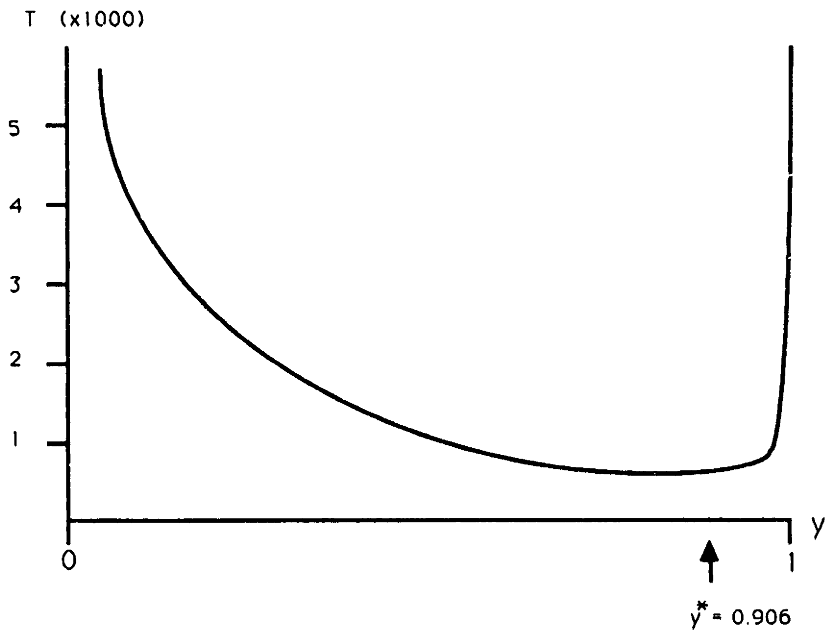

Figure 6.1: System Time vs. Travel Cost per Demand

In the uncapacitated, $m = 1$ case, we obtained a system time of the form (see Equation (3.26))

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

where $c_1, \ldots, c_3$ depended on the system parameters. This function is shown graphically in Figure 6.1 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. In general, one would select a value of y that minimizes a particular cost function g. This confirms the rather intuitive notion that there is a tradeoff between travel cost and system time in dynamic vehicle routing systems.

Similar relationships are found for the capacitated case with the exception 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, $2\lambda\bar{r}/qm$, cannot be traded-off against system time; only the local costs can.

6.3 Optimization Problems

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

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

6.3.1 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 to 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 $\lambda \bar{s} + \frac{2\lambda \bar{r}}{vq}$. 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.

6.3.2 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, one can 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.

6.3.3 Routing

Routing to minimize system time, or more generally some mixed cost, has been the main focus of this thesis. We only point out here that our results suggest a range of provably good and also quite practical policies for use at this level. Also, the fact that our analysis yields simple, closed form expressions for the asymptotic cost of these policies makes it possible to formulate and solve high-level decisions like those suggested above.

6.4 Dynamic Matching and Other Combinatorial Problems

We lastly consider applications of our heavy traffic lower bounds to problems other than dynamic vehicle routing. In particular, we shall focus on a dynamic matching problem; however, a range of combinatorial problems in the plane can be analyzed using similar ideas.

Recall that Lemmas 3.1 and 5.2 provide a lower bound on the distance to the nearest neighbor, Z^*, at the time a point is removed. Namely,

$$E[Z^*] \geq \gamma \frac{\sqrt{\Xi}}{\sqrt{N}} \tag{6.1}$$

where N is the time average number of points in queue, $\Xi = A$ for the uniform case, $\Xi = \left[\int_A f^{1/2} dx\right]^2$ for general fair policies and γ is a numerical constant which is equal to $\frac{2}{3}\frac{1}{\sqrt{2\pi}}$ in the case of Lemma 3.1 and $\frac{2}{3}\frac{1}{\sqrt{\pi}}$ in the case of Lemma 5.2. (In the case of Lemma 5.2, the inequality above is valid asymptotically as $N \to \infty$.)

In Chapter 5, we pointed out that the bound in Lemma 5.2 applies to any policy in which points are removed sequentially. (We called these policies *sequential removal rules.*) Similarly, Lemma 3.1 applies to any system in which points are removed using a nonanticipating but not necessarily sequential policy. (We call such policies *nonsequential removal rules.*) These bounds are essentially dynamic versions of the nearest neighbor bound for uniformly distributed points in the plane [36]. As such, they can be used to analyze other dynamic combinatorial problems in the Euclidean plane. We give one such example next, namely a dynamic version of Euclidean matching.

Consider the following matching problem: We are given a square of area A that receives arrivals of points according to a renewal process with intensity λ. There is cost C_1 per unit time that each point waits in the system and a cost $C_2 z$ for matching a point with a neighbor that is a distance z away. We want to match arriving points and remove them so as to minimize the time average cost

$$C = C_1 N + \lambda C_2 E[Z] = \lambda(C_1 W + C_2 E[Z]).$$

We assume points are removed in matched pairs so that policies for this problem correspond to nonsequential removal rules. The optimal cost above is denoted C^*.

Note that for any W, the lower bound (6.1) implies

$$C_1 W + C_2 E[Z] \geq C_1 W + C_2 \frac{2}{3\sqrt{2\pi}} \frac{\sqrt{A}}{\sqrt{\lambda W}}.$$

Minimizing the right hand side with respect to W we obtain the following lower bound on C^*:

$$C^* \geq (\frac{3}{2\pi})^{1/3}(\lambda C_2)^{2/3}(C_1 A)^{1/3}.$$

Now consider the following rule for matching points: divide the square into m subsquares of equal area A/m. As soon as any subregion gets two arrivals, match them and remove them. Repeat this process continually for all m subregions.

Since exactly one half of all points have no wait and the other half have a waiting time of $\frac{1}{\lambda/m} = \frac{m}{\lambda}$,

$$W = \frac{m}{2\lambda}.$$

The expected value of each matching is just the expected distance between two uniformly distributed points in a square of area A/m, which by Equation (2.8) is

$$E[Z] = a\frac{\sqrt{A}}{\sqrt{m}},$$

where $a \approx 0.52$. Combining these two expressions implies that the cost of this policy as a function of m is

$$C(m) = \frac{m}{2} + a\lambda\frac{\sqrt{A}}{\sqrt{m}}.$$

Minimizing with respect to m (ignoring integrality) we obtain

$$C(m^*) = \frac{3}{2}a^{2/3}(\lambda C_2)^{2/3}(C_1 A)^{1/3}.$$

Comparing this to the bound on the optimal value C^* above implies

$$\frac{C(m^*)}{C^*} \leq 1.24$$

so the partition policy for dynamic matching is within 25% of the optimum.

In this manner, other policies for the matching problem can be analyzed. Also, other problems such as forming dynamic MSTs, dynamic Steiner Trees, etc. can be similarly formulated. In the vehicle routing arena, dynamic versions of problems such as the Euclidean dial-a-ride problem [46] might again be analyzed using these bounds. In this problem, customers have both an origin location and a destination location and the vehicle must deliver customers from their origin to their destination. In this case, the mean nearest neighbor distance at both locations could be bounded using (6.1). These examples suggest that the lower bound (6.1) may be a useful tool for analyzing a range of dynamic combinatorial problems in the plane.

Chapter 7

Conclusions

We have examined congestion effects when operating 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, for the uniformly distributed demand case the optimal, expected system time in heavy traffic is $\Theta(\frac{\lambda A}{m^2 v^2 (1-\rho)^2})$ for uncapacitated vehicles and $\Theta(\frac{\lambda A}{m^2 v^2 (1-\rho-\frac{2\lambda \bar{r}}{mvq})^2})$ for capacitated vehicles. 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. The best of these policies is within a factor of 1.8 relative to the lower bound.

In the case of general demand distributions, we showed that there are different distributional dependencies depending on whether or not the system provides spatially fair service. We showed that the cost of providing fair service in the worst case can be arbitrarily high relative to the discriminatory case. Provable good policies where also proposed and analyzed for the general demand case. Our analysis also lead to an improvement in the constant value γ in the heavy traffic lower bound and extended the DTRP to general arrival processes.

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

These results give new insights into the problems of stability, depot location and response time under congestion for dynamic, stochastic vehicle routing systems. However, some open questions still remain in this area. A 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. Our conjecture here is that $\gamma = \beta/\sqrt{2}$, and thus the modified TSP, G/G/m and qTP policies are in fact asymptotically optimal; however, we have not been able to prove this. In Chapter 5 we saw that other combinatorial problems in the plane can be investigated using our lower bounds, and this area too seems a fruitful one for further research. 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.

Bibliography

[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 (1990), "Location of Mobile Units in a Stochastic Environment," in *Discrete Location Theory* (P.B. Mirchandani and R.L. Francis, eds.) Wiley, New York.

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

[9] Bertsimas, D. (1989), "A Vehicle Routing Problem with Stochastic Demand," to appear in *Operations Research.*

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

[11] Bertsimas, D., Jaillet, P. and Odoni, A. (1990) "A Priori Optimization," *Operations Research*, 6, 1019-1033.

[12] Bertsimas, D. and van Ryzin, G. (1990), "An Asymptotic Determination of the Minimum Spanning Tree and Minimum Matching Constants in Geometrical Probability," *Operations Research Letters*, 9, 223-231.

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

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

[17] Flores, C. (1985), "Diffusion Approximations for Computer Communications Networks," *Proceedings of Symposia in Applied Mathematics*, 31, 83-123.

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

[19] 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.

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

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

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

[23] 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.

[24] Johnson, D. (1988), talk presented at the Mathematical Progra

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

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

[27] 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.

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

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

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

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

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

[33] 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.

[34] 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.

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

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

[37] 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*)

[38] 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.

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

[40] 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.

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

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

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

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

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

[46] Stein, David M. (1978), "An Asymptotic Probabilistic Analysis of a Routing Problem," *Math. of Oper. Res.*, 3, 2, 89-101.

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

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