

STOCHASTIC AND DYNAMIC VEHICLE ROUTING IN THE EUCLIDEAN PLANE WITH MULTIPLE CAPACITATED VEHICLES

DIMITRIS J. BERTSIMAS

Massachusetts Institute of Technology, Cambridge, Massachusetts

GARRETT VAN RYZIN

Columbia University, New York, New York

(Received May 1990; revision received March 1991; accepted June 1992)

In 1991, D. J. Bertsimas and G. van Ryzin introduced and analyzed a model for stochastic and dynamic vehicle routing in which a single, uncapacitated vehicle 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. This paper extends our analysis 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. One of these policies improves by a factor of two on the best known system time for the single-server case. We then consider the case in which each vehicle can serve at most q customers before returning to a depot. We show that the stability condition in this case depends strongly on the geometry of the region. Several policies that have system times within a constant factor of the optimum in heavy traffic are proposed. Finally, we discuss extensions to mixed travel cost and system time objectives.

Vehicle routing problems are classically viewed as static and deterministic. A set of known customer locations defines an instance, and the objective is to visit customers to minimize the total travel cost subject to certain constraints (e.g., a limit on vehicle capacity). These classical problems have generated significant research interest over the years (see, for example, Lawler et al. 1985, Golden and Assad 1988) resulting in major contributions in the areas of combinatorial optimization, the analysis of heuristics and complexity theory. However, as models for the type of vehicle routing problems encountered in practice they are not always appropriate. Many real-world problems involve considerable uncertainty in the problems data. For example, locations may be known only probabilistically in advance and the demands they place on vehicle capacity may be random. In addition, requests for service often arrive sequentially in time, and again these arrival epochs may be stochastic. Finally, the objective of minimizing travel distance is not necessarily paramount; in a dynamic setting, the delivery time (wait for service) is often a more appropriate objective.

As a canonical example of an application with these characteristics, consider the following utility repair problem: A utility firm (electric, gas, water and sewer,

highway, etc.) is responsible for the maintenance of a large, geographically dispersed facilities network. The network is subject to failures which occur randomly both in time and space (location). The firm operates a fleet of repair vehicles which are dispatched from a depot to respond to failures. Routing decisions are made based on a real-time 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 the next failure. The firm would like to operate its fleet in a way that minimizes the average downtime due to failures.

There are many closely related problems to this canonical example that arise in practice. For example, consider a firm that delivers a product from a central depot to customers based on orders that arrive in real time. Orders are 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).

Other important examples are found in finished

Subject classifications Networks: stochastic, traveling salesman. Probability: stochastic model. Transportation: vehicle routing.

Area of review DISTRIBUTION, TRANSPORTATION AND LOGISTICS (SPECIAL ISSUE ON STOCHASTIC AND DYNAMIC MODELS IN TRANSPORTATION).

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. 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. In this situation, minimizing the waiting time 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.).

The complexities of such problems are often incorporated in the classical vehicle routing framework through the use of rolling horizon procedures. These involve planning routes for a fixed period into the future, often with the option of adding or deleting demands and modifying routes as time advances. See Brown and Graves (1981), Powell (1985) and Psaraftis et al. (1985) for examples of this approach. Though useful for data-intensive tactical problems, they are inherently ad hoc and hence do not give the insight necessary for strategic planning.

Several researchers have proposed alternative models that explicitly consider some combination of stochastic dynamic demands or congestion/waiting time measures. Jaillet (1988), Bertsimas (1988, 1992), and Bertsimas, Jaillet and Odoni (1990) address uncertainty in demand locations in their formulation of the probabilistic traveling salesman problem (PTSP) and the probabilistic vehicle routing problem (PVRP). Batta, Larson and Odoni (1988) and Berman et al. (1990) consider congestion effects in the context of location problems. Related work is also found in the polling system literature (Ferguson and Aminetzah 1985, Takagi 1986, Baker and Izhak 1987, Browne and Yechiali 1989) and the machine repairman literature (Stecke and Aronson 1985, Agnihothri 1988). A formulation that closely matches ours but in a network

context is the dynamic traveling salesman problem (DTSP) proposed by Psaraftis (1988). For a more detailed review of these related areas see Bertsimas and Van Ryzin (1991).

The Dynamic Traveling Repairman Problem (DTRP)

To our knowledge, however, the first comprehensive set of results on dynamic vehicle routing problems of this sort were presented in our earlier paper (1991). In it we proposed and analyzed a model called the dynamic traveling repairman problem (DTRP) that incorporates many essential features of the repair and distribution applications mentioned above. It is defined as follows: demands for service arrive according to a Poisson process with rate λ to a Euclidean service region \mathcal{A} of area A . (In this paper, we only assume that the region \mathcal{A} is bounded.) Upon arrival, demands assume an independent and uniformly distributed location in \mathcal{A} . Demands are serviced by a vehicle that travels at constant velocity v . At each location, the vehicle spends an amount of on-site service time that is generally distributed with finite first and second moments denoted by \bar{s} and \bar{s}^2 , respectively. (We let $\rho \equiv \lambda \bar{s}$ denote the fraction of time the vehicle spends in on-site service.) The objective is to find policies for routing the vehicle that minimize the average system time (waiting time plus on-site service time) of demands.

Using results from geometrical probability, queueing theory and combinatorial optimization, we were able to obtain several interesting results for this problem. In the light traffic case ($\lambda \rightarrow 0$), we showed 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

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

where X is a uniform location in \mathcal{A} . (Note that the first term above is simply the expected travel time from the median.) We extend this result to the m -vehicle case in Section 1.

In heavy traffic, we discovered a quite different and unexpected behavior. We showed that policies exist that have finite system times for all $\rho < 1$. This is surprising in that the condition is independent of the service region's size and shape; it is also the mildest stability restriction one could hope for. We then

showed that there exists a constant γ such that

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

We extend this bound to the m -vehicle case in Section 1 as well and also improve on the previous value of γ , raising it from $2/(3\sqrt{6}\pi) \approx 0.153$ to $2/(3\sqrt{2}\pi) \approx 0.266$. Note that this bound grows like $(1-\rho)^{-2}$ as $\rho \rightarrow 1$. Thus, though the stability condition is similar to that of a traditional queue, the system time increases much more rapidly as congestion increases.

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

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

where the constant γ_μ depends 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 given in Bertsimas and van Ryzin was based on forming optimal traveling salesman tours for which $\gamma_\mu = \beta \approx 0.72$, where β is the Euclidean TSP constant (see Lawler et al.). In subsection 1.3, we improve on this policy and reduce the constant to $\beta/\sqrt{2} \approx 0.51$.

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

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

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 $q \rightarrow \infty$ 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 subregions of equal size and serving each one independently using a single-server policy.

For q finite, we construct policies μ for which

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

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 (1985). 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 that 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 \bar{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. That is, $2\bar{r}/v$ is essentially the average time required to reach a set of q customers from the nearest depot (the radial cost). Dividing by q gives the average radial 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 $1-\rho-2\lambda\bar{s}/vq$, just as it does in the uncapacitated case. Note that the average radial distance \bar{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 \bar{r} (i.e., locating the depot at the median) is always optimal in heavy traffic.

The remainder of the paper is organized as follows: In Section 1 we examine the case of m vehicles without capacity constraints. We obtain lower bounds that

extend and improve on our earlier results (Bertsimas and van Ryzin). We then propose several policies for this case and analyze their performance with respect to the lower bounds. In Section 2 we examine the capacitated case. We establish a lower bound and propose and analyze two policies for the single depot, multivehicle case. As mentioned, these policies can be extended to the multidepot case provided that certain symmetry conditions hold and these are discussed in Section 2 as well. In Section 3 we discuss extending these results to minimize a mixed system time and travel cost objective. Finally, in Section 4 we give our conclusions.

1. THE M-VEHICLE, ∞ -CAPACITY DTRP

1.1. Lower Bounds

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

1.1.1. A Light Traffic Lower Bound

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

Theorem 1

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

Proof. We begin by dividing the system time of demand i T_i into three components: the waiting time of demand i due to the servers travel prior to serving i , denoted W_i^d ; the waiting time of demand i due to on-site service times of demands served prior to i , denoted W_i^s ; and 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 $i \rightarrow \infty$ gives

$$T = W^d + W^s + \bar{s}, \quad (1)$$

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

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

$$W^d \geq \frac{1}{v} \min_{|\mathcal{S}^*|=m} E \left[\min_{x_0 \in \mathcal{S}^*} \|X - x_0\| \right]. \quad (2)$$

The set of locations that achieves the minimization above is called the set of m -median locations of the region \mathcal{A} . Using the trivial bound $W_s \geq 0$ and combining with (1) and (2) establishes the theorem.

1.1.2. A Heavy Traffic Lower Bound

A lower bound useful for $\rho \rightarrow 1$ is provided by the following theorem, which generalizes and improves upon the heavy traffic bound in Bertsimas and van Ryzin.

Theorem 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 $\gamma \geq 2/(3\sqrt{2\pi}) \approx 0.266$.

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

Lemma 1

$$\bar{d} \geq \gamma \frac{\sqrt{A}}{\sqrt{N + m/2}}, \quad (3)$$

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

Before proving this lemma, observe that Theorem 2 is easy to derive from it by substituting the bound on \bar{d} into the stability condition,

$$\bar{s} + \frac{\bar{d}}{v} \leq \frac{m}{\lambda}, \quad (4)$$

which yields

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

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

Proof of Lemma 1. Consider a random “tagged” demand and define,

\mathcal{S}_0 : the set of locations of demands that 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 that arrive during the tagged demand’s waiting time ordered by their time of arrival;

X_0 \equiv the tagged demand’s location;

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

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

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

$$P\{Z_i \leq z\} \leq \frac{\pi z^2}{A}, \quad (5)$$

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

$$\bar{d} \geq E[Z^*]. \quad (6)$$

We next bound the right-hand side of (6). To do so, define an 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

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

Since $E[N_1] = N$ and $E[I_{Z_i}] = P\{Z_i \leq z\}$ is bounded according to (5), we obtain

$$P\{Z^* > z\} \geq 1 - P\{Z_0^* \leq z\} - N \frac{\pi z^2}{A}. \quad (7)$$

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

Lemma 2

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

Proof. This proof simplifies that presented in Bertsimas and van Ryzin and also extends the result to any bounded, Lebesgue (area) measurable subset of \mathcal{A} of \mathbb{R}^2 . 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 above bound to assert that

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

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

Using the result of Lemma 2 in (7) yields

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

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

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

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

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 1

$$\frac{T_{mSQM}}{T^*} \rightarrow 1 \quad \text{as } \lambda \rightarrow 0.$$

Proof. Let $j = 1, \dots, m$ 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 $\lambda_j = (A_j/A)\lambda$ denote the arrival rate to cell j and $\rho_j = \lambda_j \bar{s}$ server j 's utilization. Finally, for a uniformly distributed location $X \in \mathcal{A}$ let

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

and

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

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

$$\begin{aligned} T_{mSQM} &= \sum_{j=1}^m \frac{A_j}{A} \frac{\lambda_j (\bar{s}^2 + 4\bar{s}\bar{d}_j/v + 4\bar{d}_j^2/v^2)}{2(1 - 2\lambda_j \bar{d}_j/v - \rho_j)} \\ &\quad + \sum_{j=1}^m \frac{A_j}{A} (\bar{d}_j/v + \bar{s}), \end{aligned}$$

where the terms in the second sum are the weighted one-way travel time plus on-site service time means in each cell. As $\lambda \rightarrow 0$, $\lambda_j \rightarrow 0$ for all j and thus the contribution of the first term tends to zero, while the second term is simply

$$\frac{1}{v} E \left[\min_{x_0 \in \mathcal{D}^*} \|X - x_0\| \right] + \bar{s}$$

by construction because $\{x_j^*\} = \mathcal{D}^*$ and

$$\mathcal{A}_j = \{x \mid j = \operatorname{argmin}_k \|x - x_k^*\|\}.$$

Therefore,

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

Comparing this to Theorem 1 establishes the proposition.

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

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 new version of the TSP policy introduced in Bertsimas and van Ryzin that this dependence on m can be eliminated. In addition this policy improves on the best known constant value γ_μ for the single-server DTRP. Finally, we show that the same result holds if the service region \mathcal{A} is divided into equal-sized subregions and one of the single-server policies is applied in each region.

1.3.1. 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 (\mathbf{RA}_μ), is defined as follows.

The \mathbf{RA}_μ Policy

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

Proposition 2

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

Proof. 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

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

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

Observe that the performance guarantee for $\mathbf{RA}\mu$ has the undesirable characteristic of increasing with the number of servers.

1.3.2. A G/G/m Version of the TSP Policy

One might expect that a more intelligent allocation of customers to servers might yield a better bound. Such is indeed the case as shown by the following G/G/m version of the TSP policy. 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 k th 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 $1, \dots, n$, \mathcal{N}_2 is the set of demands $n + 1, \dots, 2n$, etc.). Assume that the server operates out of a depot at an arbitrary 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 that starts and ends at the depot. (The vehicle randomly selects one of the two possible orientations of these tours.) Optimize over n .

Proposition 3

$$\frac{T_{GGm}}{T^*} \leq \frac{m + 1}{2} \frac{\beta^2}{\gamma^2} \text{ as } \rho \rightarrow 1.$$

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 $\text{Var}[\tau] = \text{Var}[L_n]/v^2 + n\sigma_s^2$, where $\sigma_s^2 = \bar{s}^2 - \bar{s}^2$ is the variance of the on-site service time.

We next make use of a heavy traffic limit due originally to Kingman (1964) for the waiting time W in a G/G/m queue (cf. K ollerstr om 1974).

$$W \sim \frac{\sigma_a^2 + \sigma_b^2/m^2}{2\bar{t}(1 - \bar{x}/tm)} \text{ as } \frac{\bar{x}}{tm} \rightarrow 1, \tag{8}$$

where σ_a^2 and σ_b^2 are the variances for the interarrival time and service time, respectively, \bar{t} is the mean interarrival time and \bar{x} is the mean service time. Letting W_{set} denote the waiting time of a set, this limit

in our case gives

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

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

As we show below, for the policy to be stable in heavy traffic n has to be large. Thus, because the locations of points are uniform and i.i.d. in the region, we can apply the following asymptotic TSP results (cf. Steele 1981, Lawler et al. 1985):

$$\frac{E[L_n]}{n} \sim \beta \frac{\sqrt{A}}{\sqrt{n}}, \tag{11}$$

and

$$\frac{\text{Var}[L_n]}{n} \sim 0, \tag{12}$$

as $n \rightarrow \infty$. ($\beta \approx 0.72$ is the Euclidean TSP constant.) To simplify the final expressions, we have neglected the difference between $n + 1$ and n in the above expressions. (The tour includes n points plus the depot.) Since n turns out to be large, the difference is negligible. Substituting these expressions above we obtain,

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

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

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

Therefore for $\rho \rightarrow 1$, n must indeed be large, and thus using asymptotic TSP results is justified. Also, as $\rho \rightarrow 1$, $\rho + (\lambda/m)\beta_{TSP}(\sqrt{A}/(v\sqrt{n})) \rightarrow 1$ for all n satisfying (14), and thus we confirm the queue operates in heavy traffic.

The waiting time given in (13) is not the wait for service of an individual demand; it is the wait in queue for a set. The time of arrival of a set is actually the time of arrival of the last demand in that set. Therefore, to get the system time of a demand we must add to (13) the time a customer waits for its set to form, which we denote W^- , and also the time it takes to complete service of the customer once the customer's set enters service, which we denote W^+ . By conditioning on the position that a given customer takes within its set, one can show that

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

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

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

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 $\bar{\beta}$ (e.g., $\bar{\beta} = 2$ for the strip heuristic in a square (Karp 1977)). Therefore, denoting the total system time by $T_{G/m}$, for $\rho \rightarrow 1$:

$$T_{G/m} \leq \frac{\lambda(1/\lambda^2 + \sigma_a^2/m^2)}{2(1-\rho - (\lambda/m)\beta(\sqrt{A})/(v\sqrt{n}))} + \frac{n(1-m\rho)}{2\lambda} + \bar{\beta} \frac{\sqrt{nA}}{v}. \quad (15)$$

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

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

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

$$T_{G/m} \leq \frac{\lambda(1/\lambda^2 + \sigma_a^2/m^2)}{2(1-\rho)(1-y)} + \frac{\lambda\beta^2 A(1+m\rho)}{2m^2 v^2 (1-\rho)^2 y^2} + \frac{\lambda\beta\bar{\beta}A}{mv^2(1-\rho)y}. \quad (16)$$

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

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

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

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

The leading term is proportional to $\lambda A/m^2 v^2 (1-\rho)^2$.

Comparing this term to Theorem 2 establishes the proposition.

We point out that the contribution due to the queueing term (W_{ser}) 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.

1.3.3. The Modified G/G/m Policy

A modification to the G/G/m policy can eliminate the dependence on m . The analysis requires the following theorem due to Inglehart and Whitt (1970) on the behavior of the queue Σ GI/G/m (cf. Flores 1985).

Theorem 3. (Inglehart and Whitt) *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, $i = 1, 2, \dots, k$. Let $1/\mu_j$ and $\sigma_{b_j}^2$ denote the mean and variance, respectively, of the service times at server $j = 1, 2, \dots, m$. (Such a queue is denoted Σ GI/G/m.) Define $\lambda \equiv \sum_{i=1}^k \lambda_i$, $\mu \equiv \sum_{j=1}^m \mu_j$ and $\rho \equiv \lambda/\mu$. Then the mean waiting time in queue W satisfies*

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

The modified G/G/m policy 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

$$\frac{T_{MODG/m}}{T^*} \leq \frac{\beta^2}{2\gamma^2} \quad \text{as } \rho \rightarrow 1.$$

Proof. We will only sketch the proof of this proposition because the detailed analysis closely parallels that of the G/G/m policy. Observe that 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 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 the 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 Σ GI/G/m because the input process is now the superposition of k independent renewal processes, one from each of the k subregions. Using Theorem 3, one can show that W_{set} again satisfies (13). Thus, combining the three terms W^- , W^+ and W_{set} and repeating the analysis, we obtain

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

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

When $k \gg m$, the constant for this policy approaches $\beta/\sqrt{2}$. This improves on the best constant for the single-server problem ($m = 1$) of β obtained in Bertsimas and van Ryzin (essentially for the policy G/G/1). Since these constants are squared in the leading order term, the system time of this modified policy is one-half that of the previous best known policy.

1.3.4. 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.

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 $1, \dots, k$ consecutively starting at an arbitrary region. At each time $t = (\theta/k)j$, $j = 1, 2, \dots$, 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 5

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

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 (8) and proceed as in the G/G/m case.

1.3.5. 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 6

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

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

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

Comparing to the lower bound in Theorem 2 proves the proposition.

Thus, it is easy to see that any constant factor heavy traffic policy for the single-server DTRP can readily

be extended to an m -server policy with the same constant factor using independent partitions.

It is interesting to compare the policy $\mathbf{P}_{\text{MODG/G/1}}$ to the modified $\mathbf{G/G/m}$ policy. Assuming the same number of wedges k is used for both, the constant for $\mathbf{P}_{\text{MODG/G/1}}$ is $\beta\sqrt{(1 + 1/k)/2}$ while the constant for modified $\mathbf{G/G/m}$ is $\beta\sqrt{(1 + m/k)/2}$. While it is true that one can theoretically let k be an arbitrarily large constant, in practice the partitioning policy is to be preferred because it is always smaller for finite k .

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

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

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, \bar{c} is defined by $\bar{c} = \lim_{i \rightarrow \infty} E[c_i]$. 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 $r_i = \min_{x_0 \in \mathcal{D}} \|x_i - x_0\|$ and $\bar{r} = \lim_{i \rightarrow \infty} E[r_i]$. Note also that

$$\bar{r} = E \left[\min_{x_0 \in \mathcal{D}} \|X - x_0\| \right],$$

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 $\bar{c} = q\bar{d}$ without

worrying about questions of random incidence. We shall further assume that $q > 1$, because otherwise the system behaves as an ordinary $\mathbf{M/G/m}$ queue.

2.1. A Heavy Traffic Lower Bound

We begin with the following lower bound.

Theorem 4. *For $q > 1$,*

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

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 traveled 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 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

$$\begin{aligned} \bar{c} &\geq 2\bar{r} + E \left[\sum_{j=j_*+1}^{j^*} Z_j^* \right] \\ &= 2\bar{r} + E[|\Delta j|]E[Z^*], \end{aligned} \quad (17)$$

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

Lemma 3

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

Proof. 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 $\{1, 2, \dots, q\}$. By considering the joint sample space, one can show that

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

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 $E[|\Delta j|] = 1/3 (q-1/q)$. Now, if we discard outcomes with $j_1 = j_2$, which occur with probability $1/q$, the probabilities of the remaining outcomes are scaled up by a factor of $1/(1-1/q)$. Since $j_1 = j_2$ outcomes contribute nothing to $E[|\Delta j|]$ above, it therefore follows that when selecting distinct points

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

Using Lemma 3 and noting that the lower bound on $E[Z^*]$ from Lemma 1 applies in the capacitated case as well, (17) becomes

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

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

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

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

A few comments on this bound are in order. Note that if we take $q \rightarrow \infty$, the constant value is one third of the value in the uncapacitated case. This is somewhat troubling because we know that for $q = \infty$ the two problems are 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, $q > (3\lambda\bar{r}/(1-\rho))$) for which the uncapacitated bound in Theorem 2 dominates the capacitated bound in Theorem 4. Thus, when one views q as the independent parameter, Theorem 4 can often be irrelevant.

It is relevant, however, if we consider the shape of the region, 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

shows that a necessary condition for stability is

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

Thus, λ increasing toward its maximum value is equivalent to $\rho + 2\lambda\bar{r}/(mvq) \rightarrow 1$. 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.

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 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 (cf. Haimovich and Rinnooy Kan).

We conjecture that the $1/3$ value can be eliminated. Indeed, we know by Theorem 2 that for $q = \infty$ the performance guarantee is the same as in the uncapacitated case, and we show below in Proposition 9 that for $q = 2$ it is also the same. As a further motivation, one can heuristically argue that

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

as follows: A fraction $1/q$ 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

$$\begin{aligned} T^* &\geq \gamma^2 \frac{\lambda A(1-1/q)^2}{m^2 v^2 ((1-\rho-2\lambda\bar{r})/(mvq))^2} \\ &\quad - \frac{\bar{s}(1-2\rho)}{2\rho}, \end{aligned} \tag{18}$$

which, as we will show below, implies the same contrast factor bound as in the uncapacitated case for all q . Finally, note that the above is true if we restrict ourselves to optimizing over the class of policies in which the radial connections to the depot have mean \bar{r} .

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 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.

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 $\rho + 2\lambda\bar{r}/(mvq) \rightarrow 1$ (heavy traffic).

2.3.1. 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 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 .




Figure 1. Subregions of the qRP policy.

Proposition 7

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

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

Proof. 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 $E[L] = \lim_{i \rightarrow \infty} E[L_i]$. 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 $\text{Var}[L]$ by σ_L^2 , we have

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

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 (cf. Lawler et al.).

Again, the queue formed by this policy is $\Sigma GI/G/m$. Thus by invoking Theorem 3 we obtain

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

Define $\bar{\beta}$ to be the constant such that the length of the optimal tour on q uniform points in a square of area A satisfies $E[L]/\sqrt{q} = \bar{\beta}\sqrt{A}$. (If q is large, one could reasonably use the asymptotic value $\beta \approx 0.72$.) For the reader concerned about the nonsquare 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 $E[L]/\sqrt{q}$ (see Bearwood, Halton and Hammersley 1959). Since each subregion has an area A/k , substituting this expression for $E[L]/\sqrt{q}$ above gives

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

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

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

we obtain the rather complicated expression

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

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

$$y^* \approx 1 - \frac{vm}{\beta}$$

$$\sqrt{\frac{\mathcal{B}}{2A}}$$

where

$$\mathcal{B} = (1/\lambda^2 + 1/m^2(4\sigma^2/qv^2 + \sigma_L^2/qv^2 + \sigma_s^2))(1 - \rho - 2\lambda\bar{r}/(mqv))$$

For $\rho + 2\lambda\bar{r}/(mqv) \rightarrow 1$, the above approximate y^* approaches 1, thus

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

Comparing this leading term to the bound in Theorem 4 establishes the proposition.

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 $\rho + 2\lambda\bar{r}/(mqv) < 1$ for the single-depot case because T_{qRP} is finite whenever this condition is satisfied.

The existence of such a policy also allows us to establish the following theorem.

Theorem 5. *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}_{μ^*} 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 $\delta = 1 - \rho - 2\lambda(\bar{r} + \Delta\bar{r})/(mqv)$ and $\epsilon = 2\lambda\Delta\bar{r}/(mqv)$. By our qRP results and Theorem 4, if μ^* is indeed optimal, then for all $\delta > 0$, T_{μ^*} must satisfy

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

Note, however, that for $\delta \rightarrow 0$ 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.

2.3.2. The Tour Partitioning (qTP) Policy

We next analyze a policy based on the tour partitioning (TP) scheme introduced by Haimovich and Rinnooy Kan 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 $\mathcal{N}_1, \mathcal{N}_2, \dots$ 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 endpoints 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 $\mathcal{N}_2, \mathcal{N}_3, \dots$ serving sets in FCFS order. Optimize over n .

Proposition 8

$$\frac{T_{qTP}}{T^*} \leq \frac{9\beta^2(1 + m)(1 - 1/q)^2}{2\gamma^2(1 + 1/q)^2} \quad \text{as } \rho + \frac{2\lambda\bar{r}}{mqv} \rightarrow 1.$$

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 $\mathcal{N}_1 = \{X_1, \dots, X_n\}$. As shown in Haimovich and Rinnooy Kan, 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}_i) is simply

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

Removing the conditioning of \mathcal{N}_i we obtain

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

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

$$2l \frac{\bar{r}}{v} + \left(1 - \frac{l}{n}\right) \frac{E[L_n]}{v} + n\bar{s}. \quad (19)$$

The variance of the time to service a set is

$$\frac{\text{Var}[V_n]}{v^2} + n\sigma_s^2, \quad (20)$$

We shall evaluate $\text{Var}[V_n]$ shortly.

Using (19) and (20) in the $\mathbf{G}/\mathbf{G}/\mathbf{m}$ limit (8) 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,

$$W_{\text{set}} \sim \frac{\lambda(1/\lambda^2 + (1/m^2)(\text{Var}[V_n]/nv^2 + \sigma_s^2))}{2(1 - \rho - \lambda/m(2\bar{r}/vn + (1-l/n)E[L_n]/vn))}. \quad (21)$$

For $\rho + 2\lambda\bar{r}/(mvq) \rightarrow 1$ we show below that n must be large, therefore $E[L_n]/n \sim \beta\sqrt{A}/\sqrt{n}$ and $l/n \sim 1/q$. For $\text{Var}[V_n]/n$, note that

$$\begin{aligned} \text{Var}[V_n] &= \text{Var}[\tilde{L}_n] + \text{Var}[R_n] + 2\text{Cov}[\tilde{L}_n, R_n] \\ &\leq \text{Var}[\tilde{L}_n] + \text{Var}[R_n] + 2\sqrt{\text{Var}[\tilde{L}_n]\text{Var}[R_n]}. \end{aligned}$$

We can show that $\text{Var}[\tilde{L}_n]/n \sim 0$ and $\text{Var}[R_n]/n \sim \sigma_r^2/q$; therefore, dividing by n we have that $\text{Var}[V_n]/n \sim \sigma_r^2/q$. 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

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

Note that the stability condition is

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

which implies

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

Therefore assuming $n \rightarrow \infty$ as $\rho + 2\lambda\bar{r}/(mvq) \rightarrow 1$ is consistent.

Making a change of variable to

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

and using the approximation

$$y^* \approx 1 - \frac{mv}{\beta}$$

$$\sqrt{\frac{(1/\lambda^2 + \sigma_s^2/m^2 + \sigma_r^2/m^2v^2q)(1-\rho-2\lambda\bar{r}/(mvq))}{2(1-1/q)A}}$$

we finally obtain that for $\rho + 2\lambda\bar{r}/mvq \rightarrow 1$

$$T_{qTP} \sim \frac{\lambda\beta^2A(1-1/q)^2(1+m)}{2m^2v^2(1-\rho-2\lambda\bar{r}/(mvq))^2}. \quad (22)$$

Comparing this expression to the lower bound establishes the proposition.

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

The Modified Tour Partitioning (MOD q TP) 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 q TP 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 because the argument is only a slight modification of the previous proof.

Proposition 9

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

or equivalently,

$$T_{\text{MOD}_{q\text{TP}}} \sim \frac{\lambda\beta^2 A(1 - 1/q)^2}{2m^2 v^2 (1 - \rho - 2\lambda\bar{r}/(mvq))^2}.$$

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

$$\frac{T_{\text{MOD}_{q\text{TP}}}}{T^*} \leq \frac{\beta^2}{2\gamma^2},$$

which is the same as the best guarantee for the uncapacitated vehicle case. Comparing the leading behavior of $T_{\text{MOD}_{q\text{TP}}}$ above to (18), we see that the $\beta^2/2\gamma^2$ 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 \bar{r} .

Second, observe that the constant is the asymptotic value β for all values of q whereas the $q\text{RP}$ policy only achieves the asymptotic value β for large values of q . (Recall that we used an upper bound $\bar{\beta}$ on β in the $q\text{RP}$ policy.) This stems from the fact that, as the traffic intensity increases, the $q\text{RP}$ 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 $q\text{RP}$ 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 reasons, we consider the tour partitioning policy to be superior to the region partitioning policy.

2.4. Heavy Traffic Policies for Some Symmetric Multidepot Cases

We now briefly describe some multidepot 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. Furthermore, suppose these k depots induce Voronoi cells that are identical in shape and size. Then if one applies a p -vehicle policy (i.e., $q\text{RP}$ with p vehicles) in each cell, the resulting system time will be within a constant factor of the lower bound in heavy traffic. This is because 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

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

and hence the policy has a constant factor performance guarantee.

If k is large and the depots are located at k median locations, then Haimovich and Magnanti (1988) 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} , it 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 because 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 partitioned evenly among the depot locations, then it is less certain which policy is best. Indeed, there seems to be an inherent contradiction in the asymmetric case: Each set must be serviced by its closest depot to achieve a radial travel cost of \bar{r} yet the arrivals must be evenly allocated to vehicles to achieve a uniform rate of λ/m . More sophisticated bounds and/or policies are probably needed in these cases.

3. 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 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 because over an infinite time horizon the total travel distance is always infinite. Thus, rather than simply minimizing T we may 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 is easy to incorporate 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 $m(1 - \rho)/\lambda$; 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 **G/G/1** policy as defined in subsection 1.3.2) to illustrate the tradeoff. Similar results apply for the other cases.

In the uncapacitated $m = 1$ case, we obtained a system time of the form (see (16))

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

where c_1, \dots, c_3 depended on the system parameters. This function is shown graphically in Figure 2 for the case $\sigma_s^2 = 0, A = 1$ and $\bar{s} = 0.1$ 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 $\rho \rightarrow 1$, the coefficient c_2 increases much more rapidly than 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.

4. CONCLUDING REMARKS

We have examined congestion effects when operating many capacitated vehicles in a dynamic and stochastic environment. In the uncapacitated case we found that the stability condition is independent of any characteristics of the service region while in the case where each vehicle has capacity $q < \infty$, the depot location and the system geometry strongly influence the stability condition. We also showed that the distributed character of the system gives rise to behavior very different than that of traditional queues. In particular, the optimal, expected system time in heavy traffic is $\Theta(\lambda A / (m^2 v^2 (1 - \rho)^2))$ for the uncapacitated case and $\Theta(\lambda A / (m^2 v^2 (1 - \rho - 2\lambda\bar{r}/mvq)^2))$ for the capacitated case. Moreover, we found optimal policies in light traffic and several policies that have system times within a constant factor of the optimal policy in heavy traffic. We also discussed how to extend the policies to minimize a mixed objective function involving both travel and system time costs.

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. In particular, a proof of Conjecture 1 is needed to round out our understanding of the relationship between the single and the multiple-vehicle problems. A more challenging problem is to try and close the gap between the lower bound constant γ and the various policy constant 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 **G/G/m** and modified **qTP** policies are in fact asymptotically optimal; however, we have not been able to prove this. 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.

ACKNOWLEDGMENT

This research was partially supported by the National Science Foundation under grant DDM-9014751 and by grants from Draper Laboratories and the UPS Foundation.




Figure 2. System time versus travel cost per demand.

REFERENCES

- AGNIHOTRI, S. R. 1988. A Mean Value Analysis of the Traveling Repairman Problem. *IIE Trans.* **20**, 223–229.
- BAKER, J. E., AND R. IZHAK. 1987. Polling With a General-Service Order Table. *IEEE Trans. Commun.* **COM-35**, 283–288.
- BATTA, R., R. C. LARSON AND A. R. ODonI. 1988. A Single Server Priority Queueing Location Model. *Networks* **18**, 87–103.
- BEARDWOOD, J., J. HALTON AND J. HAMMERSLEY. 1959. The Shortest Path Through Many Points. *Proc. Cambridge Phil. Soc.* **55**, 299–327.
- 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.). John Wiley, New York.
- BERTSIMAS, D. 1988. Probabilistic Combinatorial Optimization Problems. Technical Report No. 194, Operations Research Center, MIT, Cambridge, Mass.
- BERTSIMAS, D. 1992. A Vehicle Routing Problem With Stochastic Demand. *Opns. Res.* **40**, 574–585
- BERTSIMAS, D., P. JAILLET AND A. ODonI. 1990. A Priori Optimization. *Opns. Res.* **38**, 1019–1033.
- BERTSIMAS, D., AND G. VAN RYZIN. 1991. A Stochastic and Dynamic Vehical Routing Problem in the Euclidean Plane. *Opns. Res.* **39**, 601–615.
- BROWN, G. G., AND G. W. GRAVES. 1981. Real-Time Dispatching of Petroleum Tank Trucks. *Mgmt. Sci.* **27**, 19–32.
- BROWNE, S., AND U. YECHIALI. 1989. Dynamic Priority Rules for Cycli-Type Queues. *Adv. Appl. Prob.* **21**, 432–450.
- FERGUSON, M., AND Y. AMINETZAH. 1985. Exact Results for Non-Symmetric Token Ring Systems. *IEEE Trans. Commun.* **COM-33**, 223–231.
- FLORES, C. 1985. Diffusion Approximations for Computer Communications Networks. *Proc. Symp. Appl. Math.* **31**, 83–123.
- GOLDEN, G., AND A. ASSAD. 1988. *Vehicle Routing: Methods and Studies*. North-Holland, Amsterdam.
- HAIMOVICH, M., AND T. L. MAGNANTI. 1988. Extremum Properties of Hexagonal Partitioning and the Uniform Distribution in Euclidean Location. *SIAM J. Disc. Math.* **1**, 50–64.
- HAIMOVICH, M., AND A. H. G. RINNOOY KAN. 1985. Bounds and Heuristics for Capacitated Routing Problems. *Math. Opns. Res.* **10**, 527–542.
- INGLEHART, D. L., AND W. WHITT. 1970. Multiple Channel Queues in Heavy Traffic, I and II. *Adv. Appl. Prob.* **2**, 150–177 and 355–364.
- JAILLET, P. 1988. A Priori Solution of a Traveling Salesman Problem in Which a Random Subset of the Customers Are Visited. *Opns. Res.* **36**, 929–936.
- KARP, R. 1977. Probabilistic Analysis of Partitioning Algorithms for the Traveling Salesman in the Plane. *Math. Opns. Res.* **2**, 209–224.
- KINGMAN, J. F. C. 1964. The Heavy Traffic Approximation in the Theory of Queues. In *Proceedings of the Symposium on Congestion Theory*, W. L. Smith and R. I. Wilkinson (eds.). University of North Carolina, Chapel Hill, 137–169.
- KÖLLERSTROM, J. 1974. Heavy Traffic Theory for Queues With Several Servers. *J. Appl. Prob.* **11**, 544–552.
- LAWLER, E. L., J. K. LENSTRA, A. H. G. RINNOOY KAN AND D. B. SHMOYS (eds.). 1985. *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*. John Wiley, Chichester, U. K.
- POWELL, W. 1985. An Operational Planning Model for the Dynamic Vehicle Allocation Problem With Uncertain Demands. Working Paper EES-85-15, Department of Civil Engineering, Princeton University, Princeton, N. J.
- PSARAFTIS, H. 1988. Dynamic Vehicle Routing Problems. In *Vehicle Routing: Methods and Studies*, B. Golden and A. Assad (eds.). North-Holland, Amsterdam.
- PSARAFTIS, H., J. ORLIN, B. BIENSTOCK AND P. THOMPSON. 1985. Analysis and Solution Algorithms of the Sealift Routing and Scheduling Problems: Final Report. Working Paper No. 1700-85, Sloan School of Management, MIT, Cambridge, Mass.
- STECKE, K. E., AND J. E. ARONSON. 1985. Review of Operator/Machine Interference Models. *Int. J. Prod. Res.* **23**, 129–151.
- STEELE, J. M. 1981. Subadditive Euclidean Functionals and Nonlinear Growth in Geometric Probability. *Ann. Prob.* **9**, 365–376.
- TAKAGI, H. 1986. *Analysis of Polling Systems*. The MIT Press, Cambridge, Mass.