The Dynamic Nearest Neighbor Policy for the Multi-Vehicle Pick-up and Delivery Problem

Patricia Kristine Sheridan†[a], **Erich Gluck**†[b], **Qi Guan**†[c], **Thomas Pickles**†[d], **Barış Balcıoğlu**‡, **Beno Benhabib**†[e]

† University of Toronto, Department of Mechanical and Industrial Engineering
5 King's College Rd., Toronto, ON M5S 3G8, CANADA,
[a]*patricia.sheridan@utoronto.ca,* [b]*erich.gluck@utoronto.ca,* [c]*aaron.guan@utoronto.ca,*
[d]*tom.pickles@utoronto.ca,* [e]*beno@mie.utoronto.ca*

‡Sabancı University,
Faculty of Engineering and Natural Sciences,
Orhanlı-Tuzla, 34956 Istanbul, Turkey,
balcioglu@sabanciuniv.edu

Abstract

In this paper, a *dynamic* nearest neighbor (DNN) policy is proposed for operating a fleet of vehicles to serve customers, who place calls in a Euclidean service area according to a Poisson process. Each vehicle serves one customer at a time, who has a distinct origin and destination independently and uniformly distributed within the service area. The new DNN policy is a refined version of the nearest neighbor (NN) policy that is well known to perform sub-optimally when the frequency of customer requests is high. The DNN policy maintains geographically closest customer-to-vehicle assignments, due to its ability to divert/re-assign vehicles that may be already en-route to pick up other customers, when another vehicle becomes available or a new customer call arrives. Two other pertinent issues addressed include: the pro-active deployment of the vehicles by anticipating in which regions of the service area future calls are more likely to arise; and, imposition of limits to avoid prohibitively long customer wait times. The paper also presents accurate approximations for all the policies compared. Extensive simulations, some of which are included herein, clearly show the DNN policy to be tangibly superior to the first-come-first-served (FCFS) and NN policies.

Keywords and Phrases: Dynamic vehicle routing; Queueing; Anticipatory behavior

1 Introduction

The vehicle routing problem (VRP) has attracted considerable attention in the literature due to its wide range of applications. It is commonly defined as seeking to serve a number of customers with a fleet of vehicles in an effective and efficient manner (Larsen, 2000). Earlier studies addressed static and deterministic versions of the VRP, in which routing decisions had to be made for a known set of demand points. Recent research, however, has focused on the dynamic versions of the VRP that consider random demand calls in time, in which routing decisions may need to be updated on-line (Ghiani et al., 2003). Berbeglia, Cordeau, and Laporte (2010) point out that in the dynamic VRP, since static solutions are not possible, researchers aim to develop solution strategies that specify which action to be taken with continuously revelaed information in time. In this paper, thus, we propose a novel dispatching policy, termed herein as the *dynamic* nearest neighbor (DNN) policy, for the dynamic pick-up and delivery problem (DPDP) under the umbrella of the VRP. In the DPDP, customers arrive in time according to a Poisson process, and each fleet vehicle serves one customer that has a distinct origin and destination, both uniformly distributed within a service region.

The study of the DPDP, through the proposed DNN policy, is relevant to many of today's practical applications, such as dispatching of taxis or limousines in city traffic, operating a trucking fleet to move full truckloads between pick-up and delivery locations, etc. The research question we explore is how to pair a vehicle from a group of available vehicles with a newly arriving customer, and how to assign a vehicle that becomes available (after delivering a customer to its destination) to another customer already waiting to be served. In the literature, (e.g., Bertsimas and Van Ryzin, 1991), the first-come-first-served (FCFS) policy, among others, has been analyzed, according to which the vehicle that becomes available is assigned to the longest awaiting customer. Lee et al. (2003) and Liao (2003) report that exploiting the technological advances in global positioning systems (GPS), cabs in Singapore are dispatched according to the nearest neighbor (NN) policy, which sends the vehicle to the customer that it can reach in the shortest time. Under the NN policy, once a customer-vehicle assignment is made, the decision is not altered. The DNN policy, we propose in this paper, permits re-routing/diverting of vehicles, en-route to other customers, if newly available information indicates better vehicle-to-customer pairings. Due to its ability to repair on-line, the DNN policy is robustly superior to FCFS and NN policies analyzed earlier in the literature, as demonstrated in this paper through a comprehensive set of simulations that consider vehicle dispatching in realistic city traffic with varying call rates during the day. We also show that the performance of the DNN policy can be further improved when vehicles are deployed in the service area by anticipating where future demand requests are more likely to arise from.

The dynamic formulation of the VRP was first discussed by Psaraftis (1988), who introduced the dynamic traveling salesman problem (DTSP) on a graph of nodes. In this problem, demands for on-site service arrive at a node according to a Poisson process. The server travels from one node to another and spends a stochastic time with a known distribution servicing the customer at the node. The objective is to maximize the expected number of customers serviced within some period of time or to minimize the expected time a customer spends in the system.

The DTSP motivated Bertsimas and Van Ryzin (1991,1993a,1993b) to model the dynamic traveling repairman problem (DTRP) for both the single and multiple servers, with or without service capacity constraints. The DTRP is the Euclidean-plane equivalent of the DTSP, i.e., the service area is not represented by a graph of nodes; locations of service requests, as in our problem, are uniformly distributed in a convex service area (or according to some other continuous distribution as in Bertsimas and Van Ryzin, 1993b). Defining the system time of a customer as the elapsed time from the instant of his/her placement of a call until his/her service is completed, Bertsimas and Van Ryzin obtain bounds on the mean system time of the "unknown" optimal policy using queueing theory and compare the performances of several policies. Among them, under light-traffic conditions, the FCFS policy, and under heavy-traffic conditions, the NN policy appear to be promising policies in minimizing the mean system time of customers, which is the performance measure we consider herein as well. Larsen, Madsen, and Solomon (2002), via simulation, demonstrate that among several policies, the NN policy is best at minimizing route distances in the partially dynamic VRP, where "partial" refers to some of the requests being known before routing decisions are made.

The problem we study is an extension to the DTRP, analyzed by Swihart and Papastavrou (1999) and Waisanen, Shah, and Dahleh (2008). The primary difference between our problem and the DTRP is that, in the DTRP, a server (e.g., the repair-person) provides service at the location of the customer for a stochastic amount of service time, whereas in our problem, a vehicle picks-up the customer and travels with him/her to another location, where the ride time is deemed to be the service time.

Advances in communication technologies since the 1990s in the form of electronic data interchange (EDI), geographic information systems (GIS), and GPS, as discussed by Psaraftis (1995), facilitate on-line dynamic routing. In this context, one may argue that the NN policy is indeed a dynamic policy since it utilizes such technologies to determine a vehicle's next customer. However, in the NN policy, an assigned vehicle cannot be re-routed (i.e., diverted) to another customer before its current assignment is accomplished. Thus, as will be detailed in Section 2 below, the primary difference between the proposed DNN and the NN policies is the ability of the former to re-route vehicles to different customers when deemed beneficial.

The benefit of diverting a vehicle that is en-route to another customer has received sparse attention in the literature. In a rare example, Regan, Mahmassani, and Jaillet (1995, 1998) and Ichoua, Gendreau, and Potvin (2000) consider the problem in the context of truck-load carriers. Both studies indicate that substantial benefits can be attained from re-routing vehicles, such as reducing the number of unserved customers and the total distance traveled when compared to not re-routing.

The proposed DNN policy performs a complete re-assignment of empty vehicles to awaiting customers each time new information is revealed and can be easily implemented via GIS, GPS, etc. We refer the reader to Cordeau and Laporte (2003), Attanasio et al. (2004), Ichoua, Gendreau, and Potvin (2000), Gendreau et al. (2006) on the use of the tabu search in the dynamic VRP and the dial-a-ride problem (DARP) (see Cordeau and Laporte, 2007, for a survey on the DARP which aims at planning a set of minimum cost vehicle routes capable of accommodating as many requests as possible from customers). While tabu search appears to be the most popular meta-heuristic, simulated annealing, genetic algorithms, the ant algorithm, and the guided local search are also widely used (Xiang, Chu, and Chen, 2008). Máhr et al. (2010) compare an agent-based solution with an online-optimization approach in a drayage problem with time windows. Parragh, Doerner, and Hartl (2008) report that comparing the performances of these heauristics directly is difficult due to the lack of standardized simulation environments. These meta-heuristics, to the best of our knowledge, have not been considered in the context of our problem, instead, the focus has been on the NN policy. We leave the implementation of meta-heuristics for future research.

In Section 2, the DNN policy is introduced as a significantly beneficial alternative to the FCFS and NN policies in minimizing the mean system time. For the simulation scenarios, it is assumed the demand rate is constant and idle vehicles remain at their last customer drop-off locations – labeled as the base-case scenarios. In Section 3, approximations are proposed for computing the mean system time under the FCFS, NN, and DNN policies. The approximation for the FCFS policy stems from the queueing literature, which scales an $M/G/1$ queue to estimate the mean system time of an $M/G/N$ queue. For the NN and DNN policies, based on the asymptotic behavior shown by Bertsimas and Van Ryzin (1993a) for such policies, we demonstrate the accuracy of simple regression models for the single vehicle, and quadratic regression models for the multi-vehicle, problems. Section 4 has a dual purpose: (i) relaxing some of the assumptions made in the previous sections by considering time-dependent customer call rates and exercising caps/limits on the maximum wait time of a customer, as well as (ii) studying the importance of being pro-active by anticipating where future customer calls may come from in order to optimally pre-position idle vehicles.

2 Problem Definition

We study a fleet of N vehicles (e.g., taxis) serving customers that randomly "appear" in a Euclidean square region $\mathcal{A}$ of area A. We assume that all vehicles travel at a constant velocity of ν. Customer requests/calls arrive in time according to a Poisson process with a rate of λ (we consider non-homogenous Poisson calls in Section 4 as a proxy for different traffic intensities observed in daily city traffic). Customers "place calls" independently, with uniformly distributed pick-up (origin) and drop-off (destination) locations. The objective at hand is to find a policy for the fleet that minimizes the mean system time of the customers in the system. The *system time of a customer* is defined herein as the elapsed time from the instant of his/her placement of a pick-up call until his/her arrival at the respective drop-off location. We denote this random variable (r.v.) in steady-state by T. The *service (ride) time of a customer*, S, on the other hand, is defined herein as the elapsed time from the instant the vehicle has picked-up the customer at his/her origin location until the instant the vehicle arrives at the customer's desired destination. As summarized here, assumptions made in this study concerning the customer arrival process and the geographical distribution of customer origins and destinations arise from the literature, e.g., Bertsimas and Van Ryzin (1991, 1993a), and Swihart and Papastavrou (1999).

As Bertsimas and Van Ryzin (1991) point out, even under an FCFS policy, the system under consideration may exhibit some differences concerning the system time from the classical $M/G/N$ queue that has the same arrival (call) process and the same service time distribution as that of S. In both systems, a customer might experience a waiting time plus its own service time that adds up to its system time. During a customer's waiting time in the $M/G/N$ queue, servers are busy serving other customers that arrived earlier in the system. In our problem, during a customer's waiting time, vehicles (servers) may not only be serving other customers, but also be driving in-between customers.

Since pre-emption of a service is not possible in any of the policies considered, service times are stochastically the same in all of them. Then, any policy, which minimizes the amount of time that vehicles spend while driving in-between customers, minimizes the mean system time of a customer. The service time of a customer is the time it takes to cover the Euclidean distance between two uniformly and independently distributed (pick-up and drop-off) points, X_1 and X_2, in a square of area A. The first two moments of the service time are, then, defined as (Larson and Odoni, 1981, p. 135):

$$E[S] = \frac{E[\|X_1 - X_2\|]}{\upsilon} = \frac{c_1\sqrt{A}}{\upsilon}, \tag{1}$$

$$E[S^2] = \frac{E\left[\|X_1 - X_2\|^2\right]}{\upsilon^2} = \frac{c_2 A}{\upsilon^2}, \tag{2}$$

where $c_1 \approx 0.52$ and $c_2 = 1/3$. Let us also define $r \equiv \lambda E[S]/N$ as the fraction of time a vehicle rides with a customer.

Bertsimas and Van Ryzin (1993a) analyze the DTRP, which is a variant of this problem. The difference of the DTRP from our problem is that a repair-person travels to the customer's location and spends a random amount of service time there, without taking the customer to a drop-off location. After the service is over, it travels to another customer if there is any waiting for service. For the DTRP, Bertsimas and Van Ryzin (1993a) present two theorems to obtain two lower-bounds for the optimal mean system time $E[T^*]$, one for the light-traffic ($\lambda \to 0$) and the other for the heavy-traffic cases ($r \to 1$). The difference that in our problem a service vehicle will be at a different location (the drop-off location of the customer) at the end of its service time, does not affect the proofs of their theorems, which render them applicable for our case as well. In particular, Theorem 2 of Bertsimas and Van Ryzin (1993a) for the heavy-traffic regime states that for some constant $\gamma \geq 2/\left(3\sqrt{6\pi}\right) \approx 0.266$,

$$E[T^*] \geq \gamma^2 \frac{\lambda A}{N^2 \upsilon^2 (1-r)^2} - \frac{E[S]\,(1-2r)}{2r}. \tag{3}$$

The fact that the optimal policy – let alone the actual value of γ – is unknown makes the assessment of the performance of a policy with respect to the optimal policy difficult. This has led us – as have earlier researchers – to compare the relative performances of alternative policies among themselves via an extensive numerical study. In this paper, we consider the FCFS, NN, and DNN policies.

While the FCFS and NN policies have been studied extensively (e.g., Bertsimas and Van Ryzin, 1991, Swihart and Papastavrou, 1999), their relative performance via numerical examples has not been assessed before in multi-server scenarios. In this paper, we not only conduct an extensive simulation study with multiple vehicles for the two policies, but also – and more importantly – propose the new DNN policy, which is shown to be superior to the NN and the FCFS policies.

2.1 The DNN Policy

A service vehicle is denoted as *assigned* when it is paired with a customer and is en-route towards his/her pick-up location. A vehicle is *busy* when it has already picked up the customer and is travelling toward the drop-off location. Otherwise, the vehicle is assumed to be *idle*. Similarly, a customer who is paired with a vehicle is denoted as an *assigned* customer, and as a *customer-in-service* when travelling in a vehicle to his/her destination. All others, who placed calls, are denoted as *unassigned* customers. One may note that, under the DNN policy, an assigned (or unassigned) customer would not know that he/she is in fact paired (or not paired yet) with a vehicle until it shows up at his/her pick-up location.

The FCFS policy is the easiest policy to implement, according to which customers are served in the order in which they arrive in the system. When a vehicle drops off its customer, it is assigned to an unassigned customer (if there is any) who has been waiting the longest. If multiple vehicles are idle, the newly arrived customer is serviced by a random vehicle.

Under the NN policy, upon dropping off a customer, the idle vehicle is assigned to the unassigned customer with the shortest distance from the vehicle's current location, if one exists. Otherwise, the vehicle remains idle. Similarly, when a new customer calls, he/she is assigned to the nearest available idle vehicle, if one exists. Otherwise, the new customer becomes an unassigned customer and must wait for service until he/she becomes the nearest one to a vehicle that becomes idle.

One notes that, in the NN policy, once a vehicle has been assigned to a customer, it cannot be re-assigned until it delivers its assigned customer to his/her destination. For instance, if a new customer who is closer to an assigned vehicle than its original customer calls, the vehicle cannot be reassigned to this new customer even if deemed beneficial. Similarly, if no unassigned customers exist and a vehicle becomes idle that is closer to an assigned customer than its originally assigned vehicle, the assignment cannot be changed, and the newly idle vehicle must remain idle until a new customer arrives. Thus, instances may arise when the re-pairing of a vehicle could lead to a lower individual system time of the closer customer, however, since this is not allowed under the NN policy, the outcome can be a potential increase in the mean system time.

The DNN policy proposed in this paper improves upon the performance of the NN policy through re-pairing. Namely, under the DNN policy, at any time, the following two conditions hold simultaneously: (i) each assigned customer is paired with the closest vehicle, provided that there is no other customer who is closer to that vehicle, and (ii) each assigned vehicle is paired with the closest waiting customer, provided that there is no other vehicle closer to that customer.

2.2 The DNN Algorithm

Since the DNN policy (as well as the NN and FCFS policies) cannot be analyzed exactly, the mean system time estimate is obtained through simulation. In the simulation model of the DNN policy, the next event could be either the call of a new customer, Event A, dropping-off a customer, Event B, or a vehicle reaching its assigned customer at the pick-up location, Event C. The simulation clock advances to the earliest of these three events, and the simulation updates the locations of vehicles and the remaining distance to travel to their assigned pick-up or drop-off locations.

When an Event C occurs, new assignments are not made; instead, the assigned vehicle becomes a busy vehicle and its assigned customer becomes a customer-in-service. For this vehicle, the remaining distance to travel is set to the distance between the customer's pick-up and drop-off locations. Events A and B, on the other hand, can lead to a possible re-assignment of some (or all) vehicles to customers as explained in the following DNN Algorithm:

If Event A occurs*: A new customer appears:*

(1) (a) *If an idle or an assigned vehicle exists, determine the closest idle or previously assigned vehicle that is strictly closer (not equidistant) to the new customer than the customer to which it is currently assigned:*

(i) *If this vehicle is an idle vehicle, proceed to Step 3;*

(ii) *If this vehicle is an assigned vehicle, unassign it from its previously assigned customer and proceed to Step 2;*

(b) *If assigned vehicles exist, but none are strictly closer to the new customer than their currently assigned customer, proceed to Step 4.*

(c) *If all vehicles are busy, proceed to Step 4.*

(2) *Assign the vehicle to the new customer. The new customer becomes an assigned customer and the previously assigned customer becomes unassigned.*

Return to Step 1 and treat the now unassigned customer as a new customer.

(3) *Assign the (idle) vehicle to the new customer. The vehicle and the new customer become an assigned vehicle and an assigned customer, respectively.*

Exit algorithm and await new trigger.

(4) *The new customer is denoted as an unassigned customer and starts waiting in a queue.*

Exit algorithm and await new trigger.

If Event B occurs: *A vehicle becomes idle:*

(5) (a) *If an assigned or unassigned customer exists, determine the closest unassigned customer or an assigned customer whose vehicle is strictly farther away (not equidistant) than the currently idle vehicle:*

(i) *If the customer is an unassigned customer, proceed to Step 7;*

(ii) *If the customer is a previously assigned customer, unassign him/her from his/her previously assigned vehicle and proceed to Step 6;*

(b) *If assigned customers exist, but none are strictly farther away from their currently assigned vehicle than the idle vehicle, proceed to Step 8.*

(c) *If there are only customers-in-service, proceed to Step 8.*

(6) *Assign the newly idle vehicle to the previously assigned customer. The vehicle becomes an assigned vehicle and the previously assigned vehicle becomes unassigned.*
Return to Step 5, and treat the now unassigned vehicle as a new idle vehicle.

(7) *Assign the (unassigned) customer to the newly idle vehicle. The vehicle and the new customer become an assigned vehicle and an assigned customer, respectively.*
Exit algorithm and await new trigger.

(8) *Keep the vehicle as idle at its current location.*
Exit algorithm and await new trigger.

2.3 The DNN Simulation Set-up

Prior to a discussion of the initial numerical results, it would also be beneficial to describe the simulation set-up herein. The simulation models were created using the *Automod* (Version 11) software package. In all the simulated experiments, the area of the square service region, A, was set to unity. A warm-up period of at most 18,000 time units ensured reaching steady-state operating conditions. Once in steady-state, in each of 20 replications, running the system for 28,800 time units resulted in sufficiently tight Confidence Intervals (CIs) around the mean system time estimates. The largest 95% CI half-width to mean system time ratios were obtained when N=1 at the highest call rate considered. This ratio equals 3.5% for the NN and DNN policies, and 12.2% for the FCFS policy. At other call rates, when N=1, the ratio is less than 1%. For N>1, 95% CI half-width to mean system time ratios do not exceed 1%, and decrease with increasing numbers of vehicles.

In the simulations, the customer call process and the pick-up and drop-off locations of the customers were synchronized by using common random numbers in corresponding replications of alternative policies in order to reduce the variation in simulation outputs, so that alternatives could be compared more accurately. Simulation models were run on a Windows-based computer with a 2.8 GHz CPU and 2GB RAM, with the DNN policy simulation taking longer to run than the FCFS and NN policies. The computation times increased with the call rate and N. When N=100, it took roughly 4.5 minutes to run each of the 20 replications of the DNN simulation model in base-case scenarios, and 12.5 minutes for the cases to be discussed in Section 4.3.

2.4 Comparison of Policies for Base-Case Scenarios – Simulation Results

In this section, via numerical examples, we present our primary comparisons of the three policies introduced in Section 2.1 for the base-case scenarios – namely, the cases in which customer call rates, and pick-up and drop-off distributions do not change over time, and service-vehicles stay idle wherever they drop-off their last customer if there are no unassigned customers present in the system. Bertsimas and Van Ryzin (1991) compare the FCFS and NN policies in the DTRP considering only a single vehicle. Via their numerical examples, they show that the NN policy outperforms the FCFS policy (and other policies they investigated) as $r = \lambda E[S]/N \to 1$, and that the FCFS policy cannot handle the higher call rates that the NN policy can.

In our numerical examples, for the FCFS policy, we consider $\rho = 2\lambda E[S]/N = 0.1$, 0.2, 0.4, 0.6, 0.8, 0.9, 0.95. Under the FCFS policy, a vehicle drives, on average, $E[S]$ time units to reach the next customer and an average of $E[S]$ time units to drive the customer to the destination. In the $M/G/N$ queue that models the FCFS policy (Section 3.1) , $2E[S]$ has to be interpreted as the service time, and ρ is the proportion of time the vehicle is busy or assigned to a customer. For the FCFS policy to be stable, one has to have $\rho<1$. The NN and DNN policies, on the other hand, yield shorter average traveling times in between customers than $E[S]$, which approaches 0 as the customer call rate increases. Thus, these two policies become unstable only when r exceeds 1. Consequently, both policies can serve higher call rates than the FCFS policy.

2.4.1 Simulation 1: Comparison of NN and FCFS Policies

In Figure 1, we present the relative performances of the NN and FCFS policies: the x-axis shows λ/N, and the y-axis shows the percent improvement of the NN policy over the FCFS policy, defined as $(E[T_{FCFS}] - E[T_{NN}])/E[T_{NN}]\,\%$ when both policies have the same number, N, of vehicles. In all the examples, the NN policy yields lower mean system time. For $N=1$, the improvement ranges from 0.16% to 343%. For $N=10(/20/100)$, the minimum improvement is 29%(/41%/68%), while the maximum improvement is 108%(/82%/75%). For $N=1$, the improvement monotonically increases with higher call rates, while for $N>1$, the improvement first declines up until $\lambda/N = 0.769$, after which the improvement starts increasing, mainly because the FCFS policy approaches the critical ρ, and its mean system time tends to infinity. Moreover, the NN policy performs better as N increases until $\lambda/N = 0.769$.

In order to study the different behavior of the improvement curves when $N=1$ and $N>1$, we estimated the mean travel time of vehicles between the last dropped-off customer and the next picked-up customer. Since the service times are stochastically the same in any policy, the better performance of the NN policy can be only attributed to lower mean travel times between two consecutive customers served. Let this travel time be denoted by the r.v., D. Under the FCFS policy, D has the same distribution as S since the vehicle's last customer's drop-off location and the next customer's pick-up location are generated independently of each other and uniformly distributed. Consequently, D has the same first two moments also given by Eqs. (1) and (2). Therefore, when $\nu = 1$, under the FCFS policy, $E[D] \approx 0.52$ and

Figure 1: The mean system time improvement of the NN policy over the FCFS policy

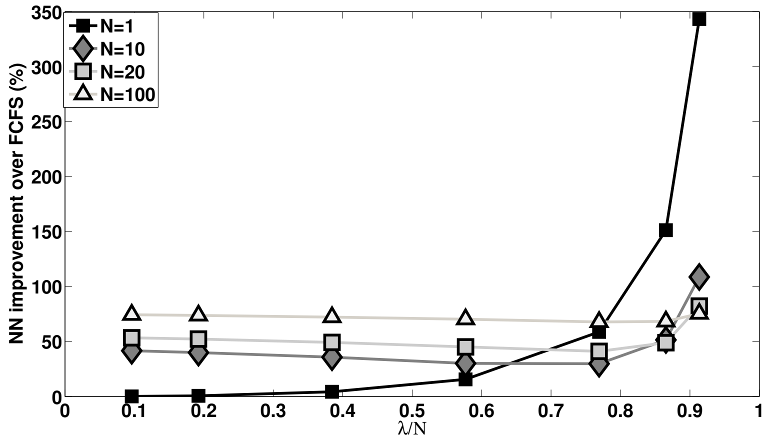

is invariant with λ/N. However, under the NN policy, Figure 2 shows how $E[D]$ changes with λ/N. As discussed above, the NN (and the DNN) policy can handle higher call rates than the FCFS policy, therefore, in addition to those ρ for the FCFS policy, we considered $\rho = 2\lambda E[S]/N = 0.99$, 1, 1.05, 1.1, 1.2, 1.3, 1.4, 1.5.

In Figure 2, we note that the larger the size of the fleet, the smaller $E[D]$ is, which explains why the NN policy performs better compared to the FCFS policy with more vehicles. However, while $E[D]$ decreases monotonically for N=1, which explains the monotonic increase in the NN policy's relative performance in Figure 1 for a single vehicle, for N>1 it increases up until a λ/N value – different for N=100 from 10 and 20 –, after which it tends to decrease. We also note that the first vertical dashed line to the left shows $\lambda/N = 0.913$ corresponding to $\rho = 0.95$, the highest call rates we considered for the FCFS policy in Figure 1, and $E[D]$ does not decrease for any N>1 we considered up until this point. We can, thus, conclude that the slight decline in the performance of the NN policy until $\lambda/N = 0.769 < 0.913$ is due to increasing $E[D]$. However, after this point the FCFS policy approaches the critical load and its mean system time increases so much that even though $E[D]$ continues to increase for a while for the NN policy, its relative performance starts improving exponentially.

Figure 2: The mean travel time between two consecutively served customers under the NN policy

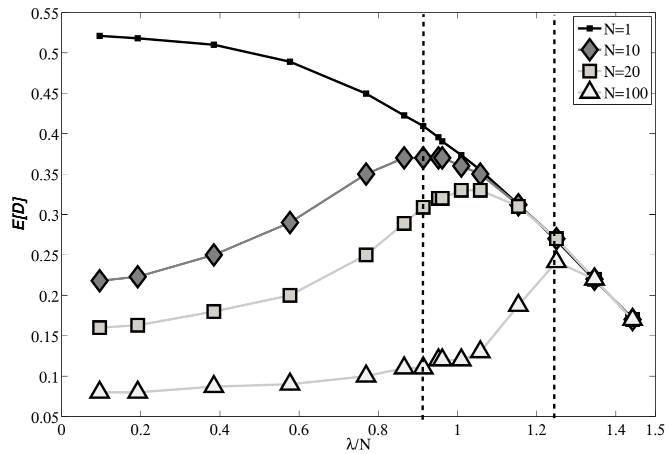

The second vertical dashed line to the left of Figure 2 is the point after which $E[D]$ becomes the same for all $N>1$ and they tend to 0. A potential explanation for the behavior of $E[D]$ for $N>1$ is that with low call rates, there are typically several idle vehicles when a new customer arrives. Since the nearest vehicle is selected for service, $E[D]$ will be less than 0.52. As λ increases there are fewer idle vehicles on average and, thus, it increases. Once λ becomes large enough, the controlling factor shifts from the number of idle vehicles to the number of unassigned customers in the system. At this point, each newly available vehicle has multiple customers to choose from, hence, $E[D]$ once again begins to decrease. It should also be noted that as N increases, the critical value of λ/N also increases.

2.4.2 Simulation 2: Comparison of DNN and NN Policies

The DNN policy is compared herein with the NN policy. In Figure 3, the x-axis shows λ/N, whereas the y-axis shows the percent improvement of the DNN policy over the NN policy defined as $(E[T_{NN}] - E[T_{DNN}])/E[T_{DNN}]\,\%$ when both policies have the same number, N, of vehicles. In all the examples, the DNN policy yields lower mean system times except when N=1, for λ/N = 1.442308. For this arrival rate the 95% CI is [-0.69%, 0.72%], and since it contains 0 both policies are deemed to be statistically not differentiable. For N=1(/10/20/100), the minimum improvement is 0.02%(/0.22%/0.11%/0.025%) while the maximum improvement is 4%(/19.64%/23.94%/25.9%).

For each N, the improvement curves are hill-shaped. As N increases, higher λ/N values are needed before the improvement curve starts inclining sharply, and similarly for the performance to peak and decline. The $E[D]$ curves for the DNN policy are similar to those in Figure 2 for the NN policy, but the values are less than or equal to those of the NN policy for the same N and λ/N values. Only when λ/N is small or large, the $E[D]$ is almost the same for both policies whereas for intermediate values, it is strictly smaller under the DNN policy. Namely, at low call rates, there are many idle vehicles when a new customer arrives and it is extremely unlikely for a vehicle rerouting to occur. As a result, there is practically no difference in the behavior of the two policies. However, as the call rate increases, the number of vehicle reassignments also increases, and at some point a difference between the policies can be seen. The greater the call rate, the more reassignments occur in the DNN policy and the average $E[D]$ decreases more than the average under the NN policy. At some point, the improvement reaches a maximum and, then, begins to decline because of the impact of a second factor. With larger call rates, there are many unassigned customers awaiting service. It is likely that any newly idle vehicle will already be near a currently unassigned customer, decreasing the chances of a beneficial vehicle rerouting to occur. As the traffic intensity continues to increase, fewer and fewer reroutes occur and the improvement declines.

3 Approximating Dispatching Policies for Base-Case Scenarios

In this section, we present analytical approximations for the three policies discussed above and test their accuracy in predicting the mean system times.

Figure 3: The mean system time improvement of the DNN policy over the NN policy

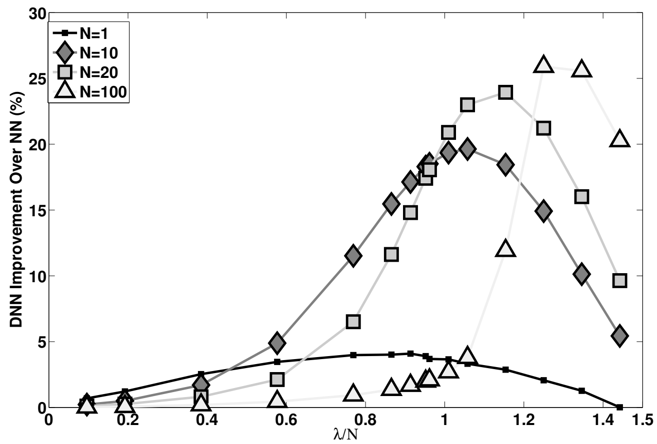

3.1 Approximating the Mean System Time Under the FCFS Policy

For the FCFS policy, the travel time between two consecutive customers, D, has the same distribution as the service time r.v. S, and its first two moments are given by Eqs. (1) and (2). Let us also define the travel time plus the actual service time a vehicle dedicates to each customer as "effective service time" (to be denoted by the r.v., $S_e = S + D$). The system can, thus, be analyzed as an $M/G/N$ queue that has Poisson calls with rate λ and S_e as its service time r.v. The effective service time, then, has a mean of $E\left[S_e\right] = 2E\left[S\right]$, and such a queue is stable if $\rho = 2\lambda E\left[S\right]/N < 1$, where the server utilization ρ corresponds to the fraction of time a vehicle is busy or assigned. Similarly, the expected system time in this queue yields the mean system time under the FCFS policy.

However, we face two problems. First, an exact solution for the mean system time in the $M/G/N$ queue does not exist. Second, the $M/G/N$ queue assumes independent and identically distributed service times. In our case, although the effective service times are identically distributed, they are not strictly independent, as pointed out by Bertsimas and Van Ryzin (1991). For the single cab in the DTRP, on-site service times are independent of travel times, and the sum of the two is the effective service time. In the DTRP, Bertsimas and Van Ryzin (1991) employ the P-K Formula for the $M/G/1$ queue to obtain the mean system time, noting that not having strict independence between consecutive effective service times does not violate the assumptions under which the P-K Formula holds (Bertsekas and Gallager, 1987, pages 142-143). On the other hand, Swihart and Papastavrou (1999) conclude that the P-K Formula can only be an approximation in the context of our problem. The second moment of the effective service time in our setting is (as in Eq. 12 of Swihart and Papastavrou, 1999)

$$E\left[S_e^2\right] = 2\left(E\left[S^2\right] + E\left[SD\right]\right).$$

Further assuming that the dependence between S and D and the dependence between consecutive effective service times can be ignored, then,

$$E\left[S_e^2\right] = 2\left(E\left[S^2\right] + E\left[S\right]^2\right),$$

where $E\left[S\right]$ and $E\left[S^2\right]$ are given in Eqs. (1) and (2), respectively.

Using $E\left[S_e^2\right]$ in the P-K Formula and adding $E\left[S_e\right]$ gives an approximation for the mean system time under the FCFS policy with a single vehicle:

$$E\left[T_{FCFS}\right] \simeq \frac{\lambda E\left[S_e^2\right]}{2\left(1-\lambda E\left[S_e\right]\right)} + E\left[S_e\right], \quad N=1. \tag{4}$$

For the multiple vehicle problems, however, we need to estimate the mean system time in the corresponding $M/G/N$ queue. One can scale an $M/G/1$ queue (see for instance Hokstad, 1978, Jagerman and Melamed, 2003) to approximate the mean queue time in the $M/G/N$ queue. Let $M/G^{(N)}/1$ queue have the same call rate λ as that of the original $M/G/N$ queue but S_e/N as the service time r.v. with $E\left[S_e\right]/N$ and $E\left[S_e^2\right]/N^2$ as its first two moments. Assuming that customers will have to wait, according to the idea of scaling, the mean waiting time in the $M/G^{(N)}/1$ queue provides an estimate for the mean waiting time in the $M/G/N$ queue. With $\rho = \lambda E\left[S_e\right]/N = 2\lambda E\left[S\right]/N$ being the probability that a customer has to wait in the $M/G^{(N)}/1$ queue, we have

$$\frac{E\left[W_{M/G^{(N)}/1}\right]}{\rho} \simeq \frac{E\left[W_{M/G/N}\right]}{P\left(W_{M/G/N}>0\right)}. \tag{5}$$

Above, since $P\left(W_{M/G/N}>0\right)$ also cannot be determined exactly, it is approximated by $P\left(W_{M/M/N}>0\right)$ of the $M/M/N$ queue with the same call rate λ and exponential service times with mean $E\left[S_e\right]$, which is known as the Erlang C function, $C\left(N,\lambda E\left[S_e\right]\right)$,

$$C\left(N,\lambda E\left[S_e\right]\right)=\frac{B\left(N,\lambda E\left[S_e\right]\right)}{1-\rho\left(1-B\left(N,\lambda E\left[S_e\right]\right)\right)},\quad B\left(N,\lambda E\left[S_e\right]\right)=\frac{\left(\lambda E\left[S_e\right]\right)^N/N!}{\sum_{j=0}^{N}\left(\lambda E\left[S_e\right]\right)^j/j!}.$$

$E\left[W_{M/G^{(N)}/1}\right]$ in Eq. (5) is determined from the P-K Formula as

$$E\left[W_{M/G^{(N)}/1}\right]=\frac{\lambda E\left[S_e^2\right]}{2\left(1-\rho\right)N^2}.$$

Using Eq. (5)

$$E\left[W_{M/G/N}\right]\simeq\frac{\lambda E\left[S_e^2\right]}{2\left(1-\rho\right)N^2}\frac{P\left(W_{M/G/N}>0\right)}{\rho}.$$

Adding $E\left[S_e\right]$ to $E\left[W_{M/G/N}\right]$ gives

$$E\left[T_{FCFS}\right]\simeq\frac{C\left(N,\lambda E\left[S_e\right]\right)}{\rho}\frac{\lambda E\left[S_e^2\right]}{2\left(1-\rho\right)N^2}+E\left[S_e\right],\quad N\geq 1. \tag{6}$$

In Table 1, we test the accuracy of the approximation given in Eqs. (4) and (6). The first column lists the ρ values considered. Under each N, there are two columns. The columns under the heading $E\left[T_{sim}\right]$ are the simulated reference values with their 95% CI half-widths given underneath them. The values under the heading $E\left[T_{FCFS}\right]$ are obtained from Eq. (4) for $N=1$ and Eq. (6) for $N>1$ with their approximation errors presented underneath them. In general, the approximation errors are quite small, at most being 4.85% for N=20 when ρ=0.95.

3.2 Approximating the Mean System Time Under the NN and DNN Policies

In line with Bertsimas and Van Ryzin (1993a), the NN and DNN policies have the same asymptotic behavior as $r=\lambda E\left[S\right]\to 1$ as

$$E\left[T_P\right]\sim\gamma_P\frac{\lambda A}{\upsilon^2(1-r)^2},$$

Table 1: Performance of the mean system time approximation for the FCFS policy

	$N = 1$		$N = 10$		$N = 20$		$N = 100$	
ρ	$E[T_{sim}]$	$E[T_{FCFS}]$	$E[T_{sim}]$	$E[T_{FCFS}]$	$E[T_{sim}]$	$E[T_{FCFS}]$	$E[T_{sim}]$	$E[T_{FCFS}]$
0.1	1.111	1.105	1.042	1.04	1.041	1.04	1.042	1.04
	±0.009	-0.54%	±0.002	-0.18%	±0.001	-0.09%	±0.001	-0.19%
0.2	1.186	1.185	1.041	1.04	1.042	1.04	1.042	1.04
	±0.007	-0.09%	±0.001	-0.09%	±0.001	-0.21%	±0.000	-0.2%
0.4	1.438	1.427	1.043	1.04	1.042	1.04	1.042	1.04
	±0.01	-0.75%	±0.001	-0.25%	±0.001	-0.19%	±0.000	-0.19%
0.6	1.924	1.911	1.06	1.055	1.045	1.042	1.042	1.04
	±0.028	-0.67%	±0.001	-0.49%	±0.001	-0.31%	±0.000	-0.22%
0.8	3.344	3.362	1.174	1.159	1.086	1.077	1.043	1.041
	±0.078	0.54%	±0.004	-1.26%	±0.001	-0.77%	±0.000	-0.25%
0.9	6.151	6.265	1.464	1.428	1.222	1.199	1.057	1.053
	0.377	1.84%	±0.01	-2.44%	±0.005	-1.81%	±0.000	-0.44%
0.95	11.839	12.069	2.088	1.999	1.554	1.479	1.109	1.099
	±1.444	1.95%	±0.057	-4.29%	±0.022	-4.85%	±0.002	-0.92%

where the coefficient γ_P depends on whether the policy P is the NN or the DNN policy. In our numerical experiments, we observed that when N=1, the following simple regression model provides an accurate fit for both policies:

$$E[T_P] \sim \gamma_P \frac{\lambda A}{\upsilon^2 (1-r)^2} + E[S_e]. \tag{7}$$

One can note that in Eq. (7), when $\lambda \to 0$, for both policies the mean system time tends to $E[S_e]$ as under the FCFS policy, since the single cab must first travel to the customer, and deliver the customer to the drop-off location; the mean of each part is $E[S]$. When the call rate is almost zero, another customer is very unlikely to appear during this time.

One can realize that the simple linear regression in Eq. (7) would not hold when $r \to 1$. However, for other r values, as will be shown, it predicts the mean system time for both policies very accurately. The advantage is that γ_P can be estimated for either policy by assuming $\nu = 1$, and one can bypass doing simulation for different velocities.

When we used the mean system time estimates for the NN and DNN policies with N=1 and $\nu = 1$, the following linear regression models were obtained, with $R^2 = 0.998$ and $R^2 = 0.996$, respectively:

$$E[T_{NN}] = 0.5103 \frac{\lambda}{(1-r)^2} + 1.04, \tag{8}$$

$$E[T_{DNN}] = 0.5039 \frac{\lambda}{(1-r)^2} + 1.04. \tag{9}$$

With γ_P values found as 0.5103 and 0.5039 as in Eqs. (8) and (9), we next assess the accuracy of Eq. (7) when N=1 but $\nu \neq 1$. Instead of presenting the results for all 15 values of $\rho = 2\lambda E[S]/N$ = 0.1, 0.2, 0.4, 0.6, 0.8, 0.9, 0.95, 0.99, 1, 1.05, 1.1, 1.2, 1.3, 1.4, 1.5, in Table 2, we present only the cases in which the maximum approximation error is observed. In Table 2, the first column lists the velocities considered; the next set of columns ρ, $E[T_{sim}]$, $E[T_{NN}]$ for the NN policy, list the ρ value with the maximum error, the estimate from simulation with its 95% CI halfwidth, and the approximated mean system time, respectively, with the approximation errors underneath them. The last three columns columns ρ, $E[T_{sim}]$, $E[T_{DNN}]$ are the same but this time for the DNN policy.

Table 2: Performance of the mean system time approximation for the NN and DNN policies when N=1

	NN			DNN		
ν	ρ	$E[T_{sim}]$	$E[T_{NN}]$	ρ	$E[T_{sim}]$	$E[T_{DNN}]$
10	1.2	0.456	0.472	1.05	0.312	0.329
		$\pm$0.002	3.62%		$\pm$0.001	5.71%
20	1.2	0.228	0.236	1.1	0.174	0.184
		$\pm$0.001	3.32%		$\pm$0.001	5.64%
100	1.2	0.046	0.047	1.1	0.035	0.037
		$\pm$0.000	3.54%		$\pm$ 0.000	5.79%

Based on the above results, we conclude that the approximation given in Eq. (7) is highly accurate. As a side note, although we are not presenting a figure, such as Figure 3, showing how the DNN policy performs with respect to the NN policy when $\nu > 1$ and N=1, we did observe that the performance of the two policies approached one another. The maximum reduction in the mean system time that the DNN policy has led to over the NN policy are 4.06 ($\pm$0.06)% when $\nu = 10$, 4.07 ($\pm$0.05)% when $\nu = 20$, and 4.08 ($\pm$0.06)% when $\nu = 100$.

When one considers the problems involving multiple vehicles, simple regression models yield poor performance. Thus, instead, by letting $K = \lambda A/(1 - r/N)^2$, the following quadratic regression model (which we also tested for N=1) gave an $R^2 > 0.99$, except for the NN policy with N=100 which resulted in $R^2 = 0.93$:

$$E[T_P] \sim \gamma_{1,P}(N)\, K^2 + \gamma_{2,P}(N)\, K + \gamma_{3,P}(N)\,. \tag{10}$$

One may note that in Eq. (10), the coefficients $\gamma_{i,P}(N)$, i=1, 2, 3, not only depend on the policy but also the number of vehicles. Table 3 lists the coefficient estimates for the NN and DNN policies.

One may observe for both policies that $\gamma_{1,P}(10)/\gamma_{1,P}(20)$ =4, $\gamma_{2,P}(10)/\gamma_{2,P}(20)$ =2, and $\gamma_{3,P}(10)/\gamma_{3,P}(20)$ =1, which suggest that the ratio of the number of vehicles and its square may appear in Eq. (10). Yet, we did not observe such ratios being preserved when the

Table 3: Coefficient estimates of the quadratic regression model for the NN and DNN policies

	NN			DNN		
N	$\gamma_{1,NN}(N)$	$\gamma_{2,NN}(N)$	$\gamma_{3,NN}(N)$	$\gamma_{1,DNN}(N)$	$\gamma_{2,DNN}(N)$	$\gamma_{3,DNN}(N)$
1	0.00024	0.464	1.09	0.0034	0.443	1.06
10	-7.5$\times 10^{-6}$	0.0068	0.756	5.2$\times 10^{-6}$	0.0037	0.728
20	-1.89$\times 10^{-6}$	0.0034	0.756	1.25$\times 10^{-6}$	0.0018	0.728
100	-4.98$\times 10^{-6}$	0.00023	0.574	-6.04$\times 10^{-6}$	5.2$\times 10^{-5}$	0.6

coefficients for N=100 or N=1 were compared to those of N=10 and N=20. We had the same problem when we considered $\upsilon \neq 1$. For N=10 and 20, we repeated our simulations assuming $\upsilon = 2,\ 5$, and 10, and validated that Eq. (10) provides accurate fits. However, for the same N value, we were unable to relate the coefficients found when different velocities were assumed. When $\upsilon \neq 1$, the relative performance of the NN and DNN policies are similar to those shown in Figure 3, and therefore, we omit to present them here.

4 Modification of the DNN Policy and its Comparison to the NN Policy in Realistic Settings

Herein, we consider three additional (realistic) settings when modifying the proposed DNN policy and comparing it to the NN policy: (i) customer call rates vary in time during the day hours, Section 4.1, (ii) traffic patterns are anticipated, Section 4.2, and (iii) customer wait times are capped, Section 4.3. Since the FCFS policy was demonstrated above to perform poorly when compared to the NN and DNN policies, it is omitted from further study.

In Section 4.1, the NN and DNN policies are analyzed via simulations under city-like conditions. The term city-like scenario refers to a model that is based on demand that peaks during rush hour traffic, as commonly experienced in metropolitan areas. In Section 4.2, anticipatory vehicle dispatching is discussed, where the modified DNN policy at hand sends idle vehicles pro-actively to regions where chances of new customer calls are higher. However, as can be expected, although mean system times can be reduced by adopting the proposed DNN policy, customers may experience excessive wait times. In Section 4.3, a capping strategy is proposed to cope with this problem by further modifying the DNN policy.

4.1 A City-Like Scenario with Varying Customer Call Rates

4.1.1 City-Like Model

In order to run a city-like scenario, the simulations considered a cyclic 24-hour time pattern, where the intervals from 10:00 am to 4:00 pm and 8:00 pm to 12:00 am are considered "normal hours", during which the base-case scenario studied in Section 2 is applied. During the period 12:00 am to 6:00 am, namely "quiet hours", the customer call rate is set to half of the call rate of the normal hours, however, the system still behaves as in the base-case scenarios in Section 2.

In order to model rush-hour customer-call patterns, during the intervals from 6:00 am to 10:00 am and 4:00 pm to 8:00 pm, the square service area was considered as two halves representing the downtown of a city and its suburbs, respectively. The former period is considered as "morning (rush) hours", during which 50% of the customers have pick-up locations in the suburbs and drop-off locations downtown. The remaining 50% of the customers have pick-up and drop-off locations without such a restriction as they were in the base-case scenarios. During the morning hours, a further "rush" occurs between 7:00 am and 9:00 am when the customer call rate is twice the call rate of the normal hours. Between 6:00 am and 7:00 am and 9:00 am and 10:00 am, the customer call rate is the same as the call rate in normal hours.

Similar to morning rush hours, the interval from 4:00 pm to 8:00 pm is considered "evening (rush) hours", during which 50% of the customers have pick-up locations downtown and drop-off locations in the suburbs. The remaining 50% of the customers have pick-up and drop-off locations without such a restriction as they were in the base-case scenarios. During the evening hours, a further "rush" occurs between 5:00 pm and 7:00 pm when the customer call rate is doubled. Between 4:00 pm and 5:00 pm and 7:00 pm and 8:00 pm, the customer call rate is the same as the call rate during normal hours.

4.1.2 Simulation Results

In simulating the city-like scenario, as before, idle vehicles under both policies stay at their drop-off locations of their last customers if there are no customers they can be assigned to. It is also assumed that $v = 1$, $A = 1$ and the call rate of the customers during normal hours are equal to the call rates λ considered in the base-case scenarios. Since the call rate is $\lambda/2$ for six quiet hours and 2λ for four rush hours each day, the mean arrival rate is $\lambda_c = 25\lambda/24$. In addition to having slightly higher call rates compared to λ in corresponding base-case scenarios, Eqs. (1) and (2) are not valid anymore due to half of customers being routed from one half of the unit square to the other for eight hours everyday. From simulations, we estimated that, in particular, the mean time a customer rides in a vehicle increases from $E[S] \approx 0.52$, for the base-case scenario, to $E[S_c] = 0.55$ for the city model. The same issue also increases the mean driving time in-between customers but its mean value depends on the customer call rate. Consequently, the mean system time under both policies in city-like traffic, that has the same λ as its normal hour call rate, increased with respect to the base-case scenarios with λ as its call rate. The increase is monotone in λ. Table 4 lists the minimum, median and maximum of these increases for both policies.

Recalling that the DNN policy loses its superiority over the NN policy with higher offered load, namely $\lambda E[S]$, we expect to see lower improvements for the DNN policy due to higher $\lambda_c E[S_c]$ when compared to those in Figure 3. In Figure 4, the x-axis shows λ/N with λ as the call rate during normal hours, whereas the y-axis shows the percent improvement of the DNN policy over the NN policy defined as $(E[T_{NN}] - E[T_{DNN}])/E[T_{DNN}]\%$ when both policies have the same number, N, of vehicles. As expected, the percentage improvements in Figure 4 are much lower than those in Figure 3, however, the same hill-shape curves recur

Table 4: Mean system time increase in city-like scenarios compared to base-case scenarios

	N	Min (%)	Median (%)	Max (%)
NN Policy	1	4.67	35.04	120.74
	10	4.74	33.05	56.66
	20	4.83	34.78	61.19
	100	4.95	42.99	68.73
DNN Policy	1	4.56	35.99	120.55
	10	4.74	38.34	58.22
	20	4.81	38.90	63.66
	100	4.96	37.25	72.39

here. With a larger fleet size, and at higher call rates, the difference between the two policies increases further, whereas as $\lambda_c E[S_c] \to 1$ or $\lambda_c E[S_c] \to 0$, the two policies tend to be identical. For N=1(/10/20/100), the minimum improvement is 0.2%(/0.22%/0.12%/0.02%) while the maximum improvement is 3.5%(/10.2%/10.4%/13.07%).

Based on our extensive studies, we conjecture that as $\lambda_c \to 0$, both policies approach the FCFS policy, for which the mean driving time is also 0.55. Then, for N=1, we expect Eq. (7) to hold if λ_c and $r_c = \lambda_c E[S_c]$ are used instead of λ and r, respectively. When we fit the following simple regression lines, they both resulted in $R^2 = 0.999$.

$$
\begin{aligned}
E[T_{NN}] &= 0.5502\frac{\lambda_c}{(1-r_c)^2} + 1.1, \\
E[T_{DNN}] &= 0.5477\frac{\lambda_c}{(1-r_c)^2} + 1.1.
\end{aligned}
$$

Similarly, redefining $K = \lambda_c/(1-r_c/N)^2$, we fit quadratic regression equations given in Eq. (10) when N>1 and all resulted in $R^2 \geq 0.99$ for the NN policy and $R^2 \geq 0.995$ for the DNN policy. However, as in the base-case examples, we were unable to relate $\gamma_{i,P}(N)$, of different N's, to one another for i=1, 2, 3.

Figure 4: The mean system time improvement of the DNN policy over the NN policy in city-like scenarios

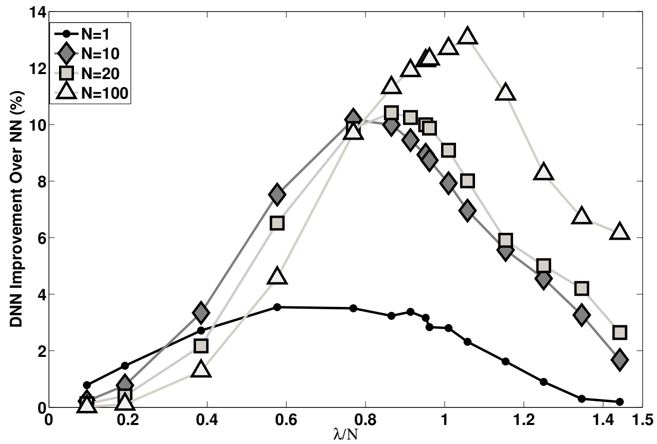

4.2 Anticipatory Behavior in City-Like Scenarios

One of the limiting assumptions made so far is that the idle vehicle remains at the drop-off location of its last customer if it has not been assigned to a new customer. In real life, however, such vehicles would usually return to some stations, or spots, to wait for new customers. However, as discussed below, if one would assume that chances of new calls would vary according to some city-like 24-hour day pattern, a dispatching policy could send the idle vehicles to wait at stations strategically in order to minimize the mean driving time to new customers.

4.2.1 Vehicle Stations

In this paper, the issue of the optimal number of vehicle stations and how to choose their optimal locations is not addressed. It is only assumed that a larger number of vehicle stations can be used for higher vehicle numbers, N, and that call patterns during rush hours need to be considered.

One can recall from Section 4.1 that during morning hours, 50% of all customers place calls in the suburbs and are dropped off downtown, and the remanning 50%, whom will be referred to as "random" customers, call equally likely from either the suburbs or downtown and get dropped off equally in either region. Consequently, if idle vehicles would remain wherever they drop off their last customers, 75% of the fleet would end up downtown. It is, thus, logical that the dispatching policy at hand would attempt to match the 75% call rate from the suburbs by having 75% of idle vehicles be present in the suburbs and not remain downtown. This can be achieved by sending back 2/3 of the idle vehicles from downtown back to the suburbs (i.e., 50% of the total fleet, which are downtown at the time) to complement the 25% of the fleet which is already in the suburbs. A similar strategy can be adopted for evening rush hours to mirror image the morning rush hours.

The proposed modified DNN policy aims at distributing vehicles uniformly among all wait stations according to the anticipated customer call patterns. Therefore, one has to have more stations in the suburbs during morning hours and fewer downtown (and vice versa during the evening hours). One may refer to the suburbs (downtown) during morning (evening) hours as the *high demand area*, and downtown (suburbs) as the *low demand area*.

4.2.2 The Anticipatory DNN Policy

Given the vehicle (wait) locations, the modified (anticipatory) DNN policy is as follows: During quiet and normal hours, all idle vehicles are directed to the closest vehicle station among those that have the fewest vehicles assigned to them. Herein, the term "assigned vehicle", refers to those vehicles waiting at the station as well as those that are en-route to the station. This implies that an idle vehicle attempts to go first to the closest vehicle station without any vehicles assigned to it. If no such station exists, it searches for the closest station with only one vehicle assigned, and so forth. This way, vehicles are evenly distributed at all the stations. The geographic distribution of the stations is discussed in detail in Appendix A.

During the morning and evening hours, on the other hand, noting that 75% of customer calls occur in the high demand area, this demand would be satisfied by keeping 75% of idle vehicles in the high demand area. This anticipatory behavior policy for morning and evening hours is adjusted as follows: During morning (evening) hours, 2/3 of the vehicles dropping off customers downtown (suburbs) and all vehicles that drop off customers in suburbs (downtown) are directed to the closest station in the suburbs (downtown) that have the least number of vehicles assigned to it. On the other hand, 1/3 of the vehicles dropping off customers downtown (suburbs) are directed to the closest station downtown (suburbs) with the least number of vehicles assigned to it. Vehicles that are at a wait station or en-route to the station are considered idle vehicles.

When a new customer calls, the modified DNN policy assigns one of the idle, or assigned, vehicles to the new customer, which may lead to re-pairing of some or all of the previous assignments. The NN policy, on the other hand, would send out whichever idle vehicle is closest to the new call. In both policies, if this turns out to be an idle vehicle en-route to a station, it changes its direction, from the wait station it was heading to, toward the customer it is assigned to.

4.2.3 Simulation Results

The call rates in Section 4.1 were used for assessing the modified (anticipatory) DNN dispatching policy and its comparison to the anticipatory NN policy. The minimum, median and maximum decreases in mean system time due to anticipatory NN/DNN dispatching compared to the city-like scenario with no anticipatory behavior are given in Table 5. For N=1, the maximum decrease for both NN and DNN policies are observed for λ = 0.096 (the normal hour call rate), which was the smallest value considered. However, for N=10, the maximum decrease is attained at λ = 0.385 and for N=20, 100 at λ = 0.577, the third and fourth smallest normal hour call rates considered. In each case, the decrease in mean system time due to anticipatory behavior diminishes as the call rate increases. This is expected since with higher offered loads, there are some unassigned customers present, and upon dropping-off a customer, vehicles are not directed to the vehicle stations, but instead to

Table 5: Mean system time decreases observed when utilizing anticipatory-behavior based dispatching

	N	Min (%)	Median (%)	Max (%)
NN Policy	1	0	0	10.23
	10	0	2.98	9.38
	20	0	3.47	9.47
	100	0.25	4.93	9.86
DNN Policy	1	0	0	9.85
	10	0	3.46	9.06
	20	0	4.66	8.94
	100	0.72	5.18	8.77

these customers; thus, the benefit of anticipatory behavior diminishes. We also note that the benefit of anticipatory behavior becomes more pronounced when there is a fleet of vehicles instead of a single vehicle.

In Figure 5, the relative performances of the NN and DNN policies are compared. For intermediate λ, for $N{>}1$, the relative performance increases by 1 to 2%. However, in general, the relative performances are similar to the results summarized in Figure 4. For N=1(/10/20/100), the minimum improvement is 0%(/0.2%/0%/0%) while the maximum improvement is 3.8%(/11.6%/12.2%/14.3%). One can, thus, deduce that the decrease in mean system time under both policies due to anticipatory dispatching is similar, which preserves the relative superiority of the DNN policy over the NN policy.

Introduction of wait stations maintains the mean riding time of the customer in vehicles as 0.55, as in city-like traffic with no anticipatory behavior in Section 4.1. However, the mean driving time between customers ($E\left[D\right]$) can be different for different customer call rates. Still assuming that as $\lambda_c \rightarrow 0$, the mean driving time is also 0.55 for N=1, we expect Eq. (7) to hold if λ_c and $r_c = \lambda_c E\left[S_c\right]$ are used instead of λ and r, respectively. When we

Figure 5: The mean system time improvement of the DNN policy over the NN policy in city-like scenarios with anticipatory behavior

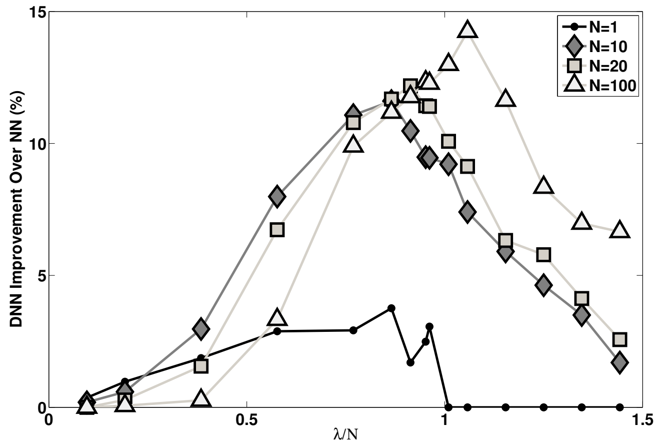

fit the following simple regression lines, they both resulted in $R^2 = 0.999$:

$$E\left[T_{NN}\right] = 0.5429\frac{\lambda_c}{\left(1-r_c\right)^2}+1.1,$$
$$E\left[T_{DNN}\right] = 0.5351\frac{\lambda_c}{\left(1-r_c\right)^2}+1.1.$$

Similarly, redefining $K = \lambda_c/\left(1 - r_c/N\right)^2$, we fit quadratic regression equations given in Eq. (10) when $N>1$ and all resulted in $R^2 \geq 0.99$ for the NN policy and $R^2 \geq 0.995$ for the DNN policy. However, as in the base-case examples or Section 4.1, we were unable to relate $\gamma_{i,P}\left(N\right)$, of different N's to one another for i=1, 2, 3.

4.3 Customer Wait-Time Limit

Although the DNN policy decreases the mean system time, it does not assure that some customers would not experience excessive wait times. Such long waiting times were observed during our simulations, since the DNN policy favors swapping longer trips for shorter trips. In order to examine how severe the cases may be, we re-examine herein the simulation output of the cases discussed in Section 4.2. For each call rate, we noted the longest wait times from each of the 20 replications and compared their average with the mean wait time estimate

Table 6: The average of longest wait times with respect to the mean wait time in city-like scenarios with anticipatory behavior under the DNN policy

N	Min	Median	Max
1	629% × 0.49 (±0.01)	1,811% × 3.14(±0.05)	3,474% × 14.41(±0.27)
10	684% × 0.17(±0.00)	1,767% × 0.72(±0.00)	4,134% × 3.86(±0.04)
20	956% × 0.13(±0.00)	1,842% × 0.62(±-0.00)	2,980% × 2.80(±0.01)
100	1,062% × 1.94 (±0.00)	2,083% × 0.39(±0.00)	3,589% × 0.10 (±0.00)

($\widehat{W}_{N,\lambda}$, for the problem with N cabs and λ as the arrival rate during normal hours). In Table 6, we present the minimum, median and maximum values obtained from this analysis as well as the corresponding $\widehat{W}_{N,\lambda}$ estimates with their 95% CIs in parenthesis. For instance, when $N = 1$, the ratio of the average of longest wait times to $\widehat{W}_{N,\lambda}$ is at its minimum when $\lambda = 0.096$ (the minimum arrival rate considered) and equals 629% of the mean waiting time which was 0.49 (±0.01). For $N = 1/10/20$, higher arrival rates tend to increase this ratio. This is not the case for $N = 100$, for which the minimum 1,062% is attained when $\lambda = 144.23$ (the highest arrival rate considered), whereas the maximum 3,589% is observed at the third smallest arrival rate $\lambda = 57.69$. The CIs indicate that although most of the observations fall close to the mean wait time estimates, the unusually long wait times can lie far from it.

4.3.1 The DNN Capping Policy

The proposed DNN policy is modified herein such that if a customer's waiting time reaches the limit, the status of this customer is prioritized. Once a change in a customer's status is noted, the closest vehicle that was assigned to another customer is redirected to this high-priority customer and serves him/her as an "absolute" assignment with no possibility of being re-paired with another customer at a later instant. This could, of course, in turn lead to the re-pairing of the remaining assigned vehicles and customers according to the DNN policy. For fair comparison, in our research, the NN policy was modified as well such that, when a high-priority customer is noted, as soon as a vehicle becomes idle, it is routed immediately to the closest high-priority customer, also as an absolute assignment.

The optimality of the actual time limit on wait times employed is beyond the scope of this study. Instead, for each problem, after estimating the standard deviation ($\widehat{\sigma}_{N,\lambda}$) and the mean wait times ($\widehat{W}_{N,\lambda}$) from the simulation output, we chose $\widehat{W}_{N,\lambda} + 6 \times \widehat{\sigma}_{N,\lambda}$ as the limit. Namely, once a customer's wait time reaches this value, it becomes a high-priority customer.

4.3.2 Simulation Results

After the capping policy was implemented, for each problem, we recorded the longest wait times from the 20 replications and compared them with those obtained when no limit was used. In Table 7, we present the reductions in longest wait times attained after implementing the capping policy. It is difficult to outline a general trend, but it may be noted that for $N > 1$ more reduction occurs under the DNN policy. This may be due to the possibility of redirecting a previously assigned vehicle to a high-priority customer as soon as it is prioritized. One may recall that under the NN policy, a high-priority customer has to wait until a vehicle drops off its customer and there are no other high-priority customers closer to this vehicle to be assigned. Additionally, we note from Table 7 that the reduction in maximum wait times due to capping tends to decrease as N increases.

The additional question that needs to be addressed is whether capping increases the mean system times as a result of reducing the longest wait times. For $N{>}1$, at $\alpha = 5\%$ significance level, the mean system times did not change most of the time for both policies. Where the change in mean system times was statistically significant, the highest increases were 0.73% ($N{=}10$, $\lambda/N = 1.009615$) for the DNN policy, and 0.67% ($N{=}10$, $\lambda/N = 0.951923$) for the NN policy. For $N{=}1$, in about half of the problems, the increase in mean system times was statistically significant at $\alpha = 5\%$. The highest increases were 5.59% and 4.13% when $\lambda = 1.346154$ for the NN and DNN policies, respectively. In summary, at the expense of no or little increase in mean system times, we can conclude that implementing caps reduced the longest wait times significantly.

Table 7: The decrease in the longest wait times after capping is implemented compared to city-like scenarios with no capping in Section 4.2

	N	Min (%)	Median (%)	Max (%)
NN Policy	1	0	113.63	225.67
	10	0	49.28	322.1
	20	0	37.16	205.4
	100	0	13.42	29.85
DNN Policy	1	0	91.98	219.37
	10	5.74	58.88	335.36
	20	18.22	52.58	227.87
	100	11.86	39.7	118.69

Given that the mean system times hardly increased, or were statistically not different from the estimates found in Section 4.2, we expect that the relative performances of the two policies will be similar to those demonstrated in Figure 5. Figure 6, where we compare the relative performances, attests to this similarity. For N=1(/10/20/100), the minimum improvement is 0%, and the maximum improvement is 4.35%(/11.5%/11.9%/14.3%).

Exercising a capping policy on maximum wait times does not change the mean riding time of the customer in the vehicle, which is 0.55, as in Sections 4.1. and 4.2. However, the mean driving time between customers, $E[D]$, depends on customer call rates. We, as in Section 4.2, assume that, as $\lambda_c \to 0$, $E[D]$ is also 0.55 for N=1, and we expect Eq. (7) to hold if λ_c and $r_c = \lambda_c E[S_c]$ are used instead of λ and r, respectively. When we fit the following simple regression lines, they both resulted in $R^2 = 0.999$.

$$E[T_{NN}] = 0.557\frac{\lambda_c}{(1-r_c)^2} + 1.1, \quad (11)$$

$$E[T_{DNN}] = 0.548\frac{\lambda_c}{(1-r_c)^2} + 1.1. \quad (12)$$

Similarly, redefining $K = \lambda_c/(1-r_c/N)^2$, we fit quadratic regression equations given in Eq. (10) when N>1 and all resulted in $R^2 \geq 0.99$ for the NN policy and $R^2 \geq 0.995$ for

Figure 6: The mean system time improvement of the DNN policy over the NN policy in city-like scenarios with anticipatory behavior and capping on wait times

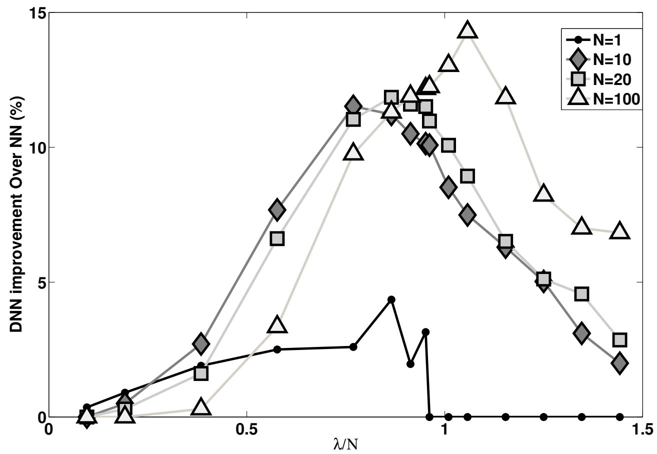

the DNN policy. However, as in cases examined in earlier sections, we were unable to relate $\gamma_{i,P}(N)$, of different N's to one another for i=1,2,3.

5 Conclusions

In this paper, a novel dynamic nearest neighbor (DNN) policy is proposed for the efficient operation of multi-vehicle fleets. The tangible advantage of the DNN-based policies over existing NN-based policies derives from its ability to divert vehicles in an on-line manner. Namely, vehicles that are already en-route to assigned customers can be diverted to others, when up-to-date real-time information makes new vehicle-customer pairings more favorable. The new policy is further augmented to cope with realistic scenarios where customer call rates and arrival distributions may vary during the day, as well as to minimize the longest customer wait times. The extensive simulation results clearly indicate the superiority of the proposed DNN policy, and its ability to account for such constraints with no or little increase in mean system times. The paper also discusses the importance of pro-active deployment of vehicles by anticipating the regions in which future calls are more likely to arise.

Approximations developed for the policies studied are also included herein. In addition to a highly accurate approximation for the FCFS policy, it is shown that simple or quadratic regression models fit almost perfectly when capturing mean system time estimates under the NN and DNN policies. As future research, assumptions made, such as the one on the customer arrival process, can be relaxed, and, furthermore, real traffic data can be used to compare these three policies. Another direction of research to follow would be implementing meta-heuristics, such as the tabu search or the simulated annealing, alongside the DNN policy, and observe which policy performs closer to the optimal policy, which may be difficult to characterize for cases with high customer arrival rates. We recommend that these research findings be put in practice by companies developing dispatching software to tangibly increase the efficiency of service-vehicle operators.

Acknowledgments

The authors thank the anonymous referees who provided invaluable suggestions to improve the manuscript.

References:

Attanasio, A., , J.-F. Cordeau, G. Ghiani, and G. Laporte. 2004. Parallel tabu search heuristics for the dynamic multi-vehicle dial-a-ride problem. *Parallel Computing*, Vol. 30, 377–387.

Berbeglia, G., J-F. Cordeau, and G. Laporte. 2010. Dynamic pickup and delivery problems. *European Journal of Operational Research,* Vol. 202, 8–15

Bertsekas, D. and R. Gallager. 1987. *Data Networks.* New Jersey: Prentice Hall.

Bertsimas, D.J. and G. Van Ryzin. 1991. A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Operations Research*, Vol. 39 (4), 601-615.

Bertsimas, D.J. and G. Van Ryzin. 1993a. Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles. *Operations Research*, Vol. 41 (1), 60-76.

Bertsimas, D.J. and G. Van Ryzin. 1993b. Stochastic and dynamic vehicle routing with general demand and interarrival time distributions. *Advances in Applied Probability*, Vol. 20, No. 4, pp. 947–978.

Cordeau, J.-F., and G. Laporte. 2003. A tabu search heuristic for the static multi-vehicle dial-a-ride problem. *Transportation Research Part B*, Vol. 37, 579–594.

Cordeau, J.-F. and G. Laporte. 2007. The dial-a-ride problem: Models and algorithms. *Annals of Operations Research*, Vol. 153, 29–46.

Hokstad, P. 1978. Approximation for the M/G/m queue. *Operations Research,* Vol. 26 (3), 510–523.

Ichoua, S., M. Gendreau, and J.-Y. Potvin. 2000. Diversion issues in real-time vehicle dispatching. *Transportation Science*, Vol. 34, 426–438.

Gendreau, M., F. Guertin, J-Y. Potvin, and R. Séguin. 2006. Neighbor search heuristics for a dynamic vehicle dispatching problem with pick-ups and deliveries. *Transportation Research Part C*, Vol. 14, 157-174.

Ghiani, G., F. Guerriero, G. Laporte, and R. Musmanno. 2003. Real-time vehicle routing: solution concepts, algorithms and parallel strategies. *European Journal of Operational Research*, Vol. 151, 1-11.

Jagerman, D. L. and B. Melamed. 2003. Models and approximations for call center design. *Methodology and Computing in Applied Probability*, Vol. 5 (2), 159-181.

Larsen, A. 2000. *The Dynamic Vehicle Routing Problem.* Ph.D.-thesis. IMM-PHD-73, IMM.

Larsen, A., O. Madsen, and M. Solomon. 2002. Partially dynamic vehicle routing-models and algorithms. *Journal of the Operational Research Society*, Vol. 53, 637-644.

Larson, R. and A. Odoni. 1981. *Urban Operations Research.* New Jersey: Prentice Hall.

Lee, D. H., H. Wang, R. L. Cheu, and S. H. Teo. 2003. A taxi dispatch system based on current demands and real-time traffic conditions. *82nd Transportation Research Board Annual Meeting of the Transportation Research Board*, in CD-ROM, 12 - 16 Jan, Washington DC, USA.

Liao, Z. 2003. Real Time Taxi Dispatching Using Global Positioning Systems. *Communications of the ACM.* Vol. 46, No. 5, 81-83.

Máhr, T., J. Srour, M. de Weerdt, and R. Zuidwijk. 2010. Can agents measure up? A comparative study of an agent-based and on-line optimization approach for a drayage problem with uncertainty. *Transportation Research: Part C*, Vol. 18, No. 1, 99-119.

Parragh, S. N., K. F. Doerner, and R. F. Hartl. 2008. A survey on pickup and delivery problems. *Journal für Betriebswirtschaft*, Vol. 58, No. 2, 81-117.

Psaraftis, H. N. 1988. Dynamic Vehicle Routing, in *Vehicle Routing: Methods and Studies.* B. L. Golden and A. A. Assad (eds), 223-248, North-Holland, Amsterdam

Psaraftis, H. N. 1995. Dynamic vehicle routing: Status and prospects. *Annals of Operations Research*, Vol. 61, 143-164.

Regan, A. C., H. S. Mahmassani, and P. Jaillet. 1995. Dynamic decision making for commercial fleet operations using real-time information. *Transportation Research Record*, No. 1537, 91-97.

Regan, A. C., H. S. Mahmassani, and P. Jaillet. 1998. Evaluation of dynamic fleet management systems. *Transportation Research Record*, No. 1645, 176-184.

Swihart, M. R. and J. D. Papastavrou. 1999. A stochastic and dynamic model for the single vehicle pick-up and delivery problem. *European Journal of Operations Research*, Vol. 114 (3) 447-464.

Waisanen, H.A., D. Shah, and M.A. Dahleh. 2008. A dynamic pickup and delivery problem in mobile networks under information constraints. *IEEE Transactions on Automatic Control*, Vol. 53, 1419-1433.

Xiang, Z., C. Chu, and H. Chen. 2008. The study of a dynamic dial-a-ride problem under time-dependent and stochastic environments. *European Journal of Operational Research*,

Table A.1: Number of Vehicle Stations

	Number of Vehicle Stations (N_V, N_H)		
N	Low Demand Area	Uniform Demand Area	High Demand Area
1	2 (2,1)	1 (1,1)	2 (2,1)
10	6 (2,3)	9 (3,3)	16 (4,4)
20	12 (4,3)	20 (4,5)	30 (6,5)
100	48 (6,8)	100 (10,10)	144 (12,12)

Vol. 185, 534-551.

Appendix A Vehicle Stations

Herein, the vehicle (wait) stations are distributed uniformly over the service region according to the number of vehicles within the fleet. Due to the square-shaped region, each station is located at the intersection of fictitious vertical and horizontal lines. In Table A.1, the numbers of stations are given for the different numbers of vehicles considered in our simulations, N. The number of vertical (N_V) and horizontal (N_H) lines are also given (in parenthesis) in this table. When N_V and N_H are not equal to the square root of the number of vehicle stations, we choose them to be the divisors of the number of locations which were closest to one another. In such cases, it does not matter which one is N_V or N_H, when there is uniform demand, since the service area is a square. One may recall from Section 4.1 that during quiet and normal hours, the system experiences uniform demand. During morning and evening hours, one half of the service area is the high- and the other side is the low-demand area. During these times, half the numbers of vehicle stations listed under the low-demand area and high-demand area are located in their respective halves, as the numbers provided in Table A.1 are for a full unit square.

In order to determine the coordinates of the service stations, let us assume that the (x,y) coordinates of the lower left corner of the service area are at the origin, (0,0). The

Figure A.1: Locations of vehicle stations when N=10 and demand is uniform during quiet and normal hours

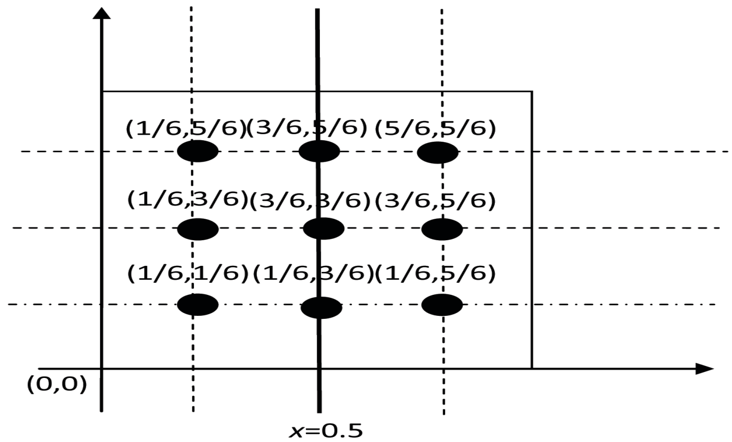

vertical lines are, then, at $x = 1/2N_V + i/N_V$, i=0,...,N_V-1 and horizontal lines are at $y = 1/2N_H + i/N_H$, i=0,...,N_H-1. The example distribution in Figure A.1 corresponds to the case of N=10 under uniform demand, such as during quiet and normal hours, for which the 9 vehicle stations are noted by the black dots. The bold vertical line at x=0.5 represents the boundary between downtown and the suburbs.

For the anticipatory NN/DNN policies only half of the number of vehicle stations presented in Table A1, under the columns low- and high-demand areas, respectively, are actually used. For the low-demand area, in Table A.1, we note 6 possible vehicle stations, however, we only use 3. Similarly, for the high-demand area, we note 16 vehicle stations, however, we only use 8. The position of these vehicle stations is determined by which half of the vehicle stations fall on the appropriate side of the dividing line between the suburbs and downtown. If one assumes that, in Figure A.2, the low-demand area on the left is downtown, this example would be representing the vehicle stations for morning hours. During evening hours, the locations will be the mirror image of these ones with respect to x=0.5.

Figure A.2: Locations of vehicle stations during morning and evening hours when N=10