DISPATCHING OF MULTIPLE SERVICE VEHICLES IN THE DYNAMIC-DIAL-A-RIDE PROBLEM

Thomas Pickles[(a)], Barış Balcıoğlu[(b)], Beno Benhabib[(c)]

[(a) (b) (c)] University of Toronto
Department of Mechanical & Industrial Engineering
5 King's College Road Toronto ON M5S 3G8 Canada

[(a)] tom.pickles@utoronto.ca, [(b)] baris@mie.utoronto.ca, [(c)] beno@mie.utoronto.ca

ABSTRACT
Dispatching of multiple service vehicles is studied for the Dynamic Dial-A-Ride Problem (Dynamic DARP). In the Dynamic DARP service vehicles transport demands from pick-up to drop-off locations which are both independently and uniformly distributed within a unit square service region. Three policies are compared: First Come First Serve (FCFS), Nearest Neighbour (NN), and Dynamic Nearest Neighbour (DNN). Simulation results show NN outperforms FCFS by margins of up to 40%, and DNN provides further improvements up to 10% over NN. Analytical approximations are then developed for the multiple vehicle FCFS policy and the single vehicle varying velocity NN policy. The service environment is then expanded to more realistic city-like conditions. Simulation results confirm the relative policy performance holds. Finally, anticipatory vehicle routing is applied such that idle service vehicles are proactively routed. Anticipatory routing provides up to an 18% improvement over the reactive policies.

Keywords: Dynamic Vehicle Routing, Policy Comparisons, Approximations, Simulation

1. INTRODUCTION

The Vehicle Routing Problem (VRP) is defined in simple terms as seeking to serve a number of customers with a fleet of vehicles in an effective and efficient manner. Due to its practicality and wide range of applications the VRP has attracted considerable academic attention. Traditionally, studies have focused on static and deterministic versions of the problem using a set of predetermined demand locations (Xu, 1994). However, in reality, demands (customers or objects) often arrive randomly in time and therefore require continuous dispatching processes (Bertsimas and Van Ryzin, 1991). In most routing situations the problem is inherently dynamic and stochastic in nature. Furthermore, the classical objective of VRPs, to minimize travel distance and associated direct travel costs, may not always prove to be the most important factor. In many dynamic situations, such as the ones mentioned above, minimizing the wait for service time is more important than minimizing travel cost.

Bertsimas and Van Ryzin (1991) were the first to study a dynamic and stochastic version of the VRP. The authors introduced the Dynamic Travelling Repairman Problem (DTRP) with the objective of minimizing the wait for service. The DTRP is defined as follows: demands arrive in time according to a Poisson process in a uniformly distributed random location within a Euclidean region. The demands are then serviced at that location for an independent and identically distributed amount of time by a single repair vehicle which travels to the demands with constant unit velocity. The authors analyzed several solution policies: First Come First Serve (FCFS), Partitioning (PART), Travelling Salesman Policy (TSP), Space Filling Curve (SPC), and Nearest Neighbour (NN). Of particular interest for this paper are the FCFS and NN policies. FCFS is the simplest policy where demands are serviced in the order in which they arrive. In the NN policy, the demand closest to the service vehicle is serviced first, independent of the order in which the demand arrived. For the single server DTRP under heavy traffic conditions the NN policy was shown to perform with the lowest average system time. Bertsimas and Van Ryzin (1993) extended the DTRP to the multiple service vehicle case but were primarily concerned with developing analytical system bounds and thus did not conduct further trials of the NN policy. Ozesenli and Demirel (2005) used simulation for the DTRP in a more realistic city-like environment in which the repairman will not have a constant velocity and proposed the Shortest Arrival Time (SAT) policy.

An extension to the DTRP is the Dynamic Pick-up and Delivery Problem (DPDP) where the service vehicle must transport each demand between an origin and a destination. The primary difference between the DPDP and the DTRP is that in the DTRP the vehicle spends time at the location of each demand to serve it and thus does not change location during service, but in the DPDP the vehicle transports the demand and thus does change location during service.

The one-to-one DPDP can be further subdivided into the Dynamic Stacker Crane Problem (Dynamic SCP), Dynamic Vehicle Routing Problem with Pick-up and Deliveries (Dynamic VRPPD), and the Dynamic Dial-A-Ride Problem (Dynamic DARP) (Berbeglia, Cordeau and Laporte, 2010). The Dynamic SCP deals with optimizing a trucking fleet to move full truck loads directly from its pick-up to delivery location. The Dynamic VRPPD applies when vehicles can serve more than one request at the same time and is normally applied to courier services. Of particular interest to this paper is the Dynamic DARP where requests consist of users that need to be transported from a pick-up location to a drop-off location. Typical applications of Dynamic DARP include taxi services in cities (Berbeglia, Cordeau and Laporte, 2010).

The Dynamic DARP was introduced by Swihart and Papastavrou (1999) with the objective to minimize the expected time in the system for demands. The problem definition is analogous to that of Bertsimas and Van Ryzin (1991) with the alteration that each demand must be transported to independent and uniformly distributed delivery locations, which are independent of the pick-up locations. Three policies were compared: Sectoring, Stacker Crane, and Nearest Neighbour (NN). As in the DTRP, the NN policy delivered the lowest system times under heavy traffic conditions for the Dynamic DARP. Xiang, Chu, and Chen (2008) introduced several constraints, such as break down of vehicles, scheduled and dynamic arrivals, maximum work time for each driver, etc., in the DARP problem in which demand and delivery points are the vertices of a network.

This paper explores various policies for the Dynamic DARP with the primary objective being to minimize system time. In particular, various policies are first explored and compared in a simple base case scenario. Next, analytical expressions are derived for certain policies such that full simulations need not be run to obtain ballpark results. The policies are then applied to a more realistic city-like region to test if the relative performance holds under such conditions. Finally, anticipatory behaviour is explored to examine if its use results in improved performance.

2. POLICIES

The three policies studied in this paper are now described in detail.

2.1. First Come First Serve (FCFS)

In FCFS demands are simply served in the order in which they arrive in the system. If multiple vehicles are idle, the demand is serviced by a random vehicle. Idle service vehicles remain in their current location.

2.2. Nearest Neighbour (NN)

In NN the demand closest to the vehicle is serviced first, independent of the arrival order. Specifically, if there is more than one customer waiting a vehicle that drops off a demand will next service the demand with the pick-up location that is the smallest distance from the service vehicle's current location. If a new demand arrives and multiple service vehicles are idle, the demand is serviced by the nearest vehicle. Conversely, if there are no idle vehicles the demand must wait for service until it is the nearest demand to a newly idle vehicle. Once a service vehicle has been assigned to a demand it cannot be rerouted until the demand reaches its drop-off location.

2.3. Dynamic Nearest Neighbour (DNN)

The DNN policy is similar to NN with some minor adjustments. Again, demands are serviced by the nearest vehicle, independent of the arrival order. However, unlike NN, under certain circumstances service vehicles can be rerouted to a new demand without completing service for the existing demand.

Before entering a more detailed explanation, the necessary terminology is introduced. A service vehicle is said to be 'assigned' when it is travelling towards a demand's pick-up location. A vehicle is 'busy' when it has picked up the demand and is travelling between the pick-up and drop-off locations. If a new demand arrives and there are no idle service vehicles, the demand may be serviced by an 'assigned' vehicle if the new demand is closer to the 'assigned' vehicle than the currently assigned demand. If the new demand satisfies this criterion for more than one 'assigned' service vehicle, the new demand is serviced by the vehicle which is nearest to it. The service vehicle is re-assigned from the current demand to the new one, and changes its travelling direction towards the new demand. The dropped demand remains where it is and is now treated as any other waiting demand. This process is referred to as a 'vehicle assignment reroute'. However, if when the new demand arrives there are one or more idle vehicles the new demand is always serviced by the nearest idle vehicle regardless of the location of any 'assigned' vehicles. As in NN, when a service vehicle becomes newly idle it is 'assigned' to the nearest waiting demand. It is important to mention that despite its name the DNN is only a partially dynamic policy

3. COMPARISON OF BASE CASE POLICIES

As a basis for discussion, an understanding of the relative performance of the three policies in the basic environment as defined by Swihart and Papastavrou (Swihart and Papastavrou 1999) is required. This chapter summarizes the performance of the FCFS, NN, and DNN policies, but first begins with a detailed problem definition of the base case.

3.1. Problem Definition

The service vehicles are located in a square region of area A=1. The demands for service arrive randomly in time according to an exponential distribution with arrival rate λ. The expected time between arrivals is defined as $x=1/\lambda$. The demand pick-up locations are independently and uniformly distributed within the square region. The drop-off locations are also independently and uniformly distributed and are independent of the pick-up locations. The distances between locations are defined in the Euclidean plane. In the base case the service vehicles travel with a constant unit velocity and they travel in a straight path between locations (i.e., there are no roads). Under these conditions, distance and time are equivalent. When a service vehicle is idle it remains in its current location. The number of service vehicles is denoted by N, where N=1 for the single vehicle case and N>1 for the multiple vehicle cases.

The service time, s, is defined from when the service vehicle begins travelling to the demand's pick-up location until the demand reaches its drop-off location. The service time consists of two components: s_w, the time the service vehicle spends travelling to the demand's pick-up location, and s_t, the time the service vehicle spends transporting the demand between the pick-up and drop-off location. Thus, $s = s_t + s_w$. The rate of services is defined as $\mu = 1/s$. Wait time, W, is equal to the time from when the demand arrives until the service vehicle begins travelling towards its pick-up location. The actual wait time of the demand until the service vehicle arrives is expressed as $W + s_w$. The total system time, T, is defined from the instant the demand arrives until it reaches its drop-off location. Therefore, $T = s + W$. The overall objective is to minimize the system time, T.

The utilization, ρ, of an individual service vehicle is defined as the proportion of the time the vehicle is servicing a demand relative to the overall time. Utilization under the FCFS policy is related to the arrival and service rates by

$$\rho = \frac{\lambda E[s]}{N} = \frac{\lambda}{\mu N} \quad (1)$$

For the system to remain stable, ρ must remain less than one. The utilization is closely related to the traffic intensity (arrival rate) of the system (Xu, 1994). In general, the greater the traffic intensity the larger the value of ρ.

3.2. Results

The simulation results comparing the system times of the FCFS, NN, and DNN policies for the cases of 1, 10, 20, and 100 service vehicles are considered. Under all four cases the NN and DNN policies significantly outperform the FCFS policy. Furthermore, DNN shows a small improvement over NN. For all three policies the system becomes more efficient as the number of vehicles increases. The relative performance of the policies is examined further, beginning with a comparison between NN and FCFS, followed by DNN and NN.

3.2.1. Comparison of NN to FCFS

The NN policy outperforms the FCFS policy by different magnitudes depending on the arrival rate and the number of service vehicles. The relative performance of the two policies is compared in Figure 1.

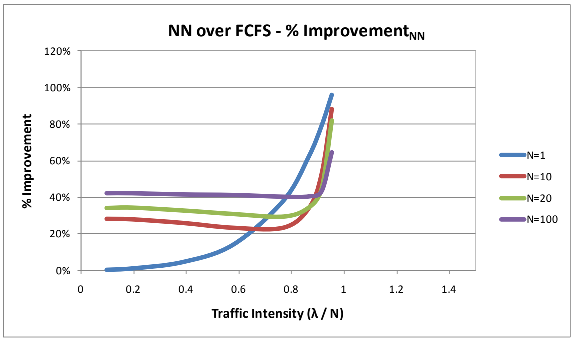

Figure 1: System Time Percent Improvement for NN Over FCFS

The typical improvement of NN over FCFS ranges from around 25% to 40%, but it increases exponentially near the FCFS traffic intensity limit. At lower arrival rates the improvement is greater for a larger number of vehicles. However, in the single vehicle case the improvement of NN over FCFS continually increases as λ increases while in the multiple vehicle cases it is fairly steady until asymptotic behavior. Furthermore, it appears that as N increases the relative performance becomes less affected by the arrival rate.

It is also of interest to examine the variation in the service time, s, for the NN policy as the number of vehicles changes. The results are summarized in Figure 2.

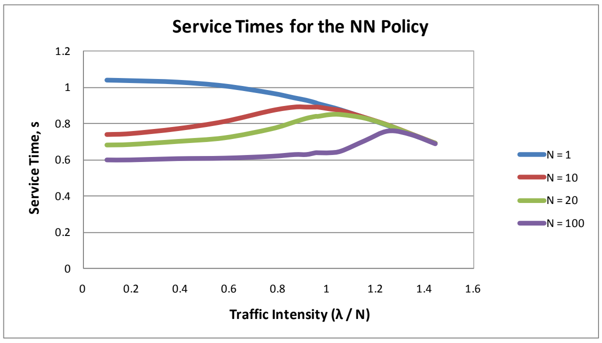

Figure 2: Service Times for NN Policy

The single vehicle case behaves as one would expect. As λ increases the service time decreases. With more customers in the system to choose from, the service vehicle finds demands that are, on average, nearer to the vehicle's current location thus decreasing s_w and therefore s. However, for the multiple vehicle case the results are more interesting. As λ increases, s also increases until some critical point where s then behaves as in the single vehicle case. A potential explanation for this phenomenon is that with low arrival rates there are typically several idle vehicles when each demand arrives. Since the nearest vehicle is selected for service, s_w will be less than 0.52 (mean distance between two points independently generated uniformly distributed points in the unit square). As λ increases there are fewer idle vehicles on average and thus s_w increases. Once λ becomes large enough, the controlling factor shifts from the number of idle vehicles to the number of demands in the system. At this point each newly available vehicle has multiple demands to choose from so s_w once again begins to decrease. It should also be noted that as N increases the critical value of λ also increases.

3.2.2. Comparison of DNN to NN

In order to understand the incremental gains of DNN over NN it is logical to compare DNN directly to NN rather than to FCFS. This comparison is summarized in Figure 3.

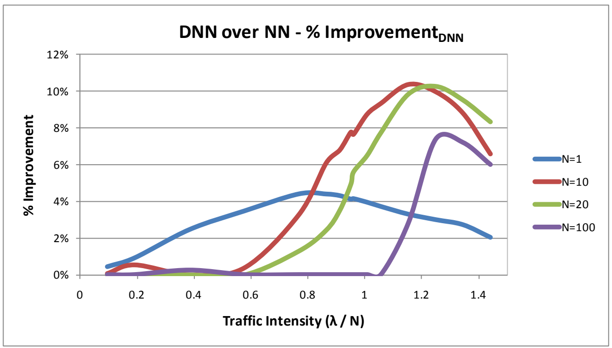

Figure 3: System Time Percent Improvement for DNN Over NN

The maximum improvement of DNN over NN ranges from around 4% to 10% depending on the number of service vehicles. With a single service vehicle the maximum improvement is a modest 4%. With 10 or 20 vehicles the improvement reaches a maximum of just over 10% and in the 100 vehicle case the improvement retreats to just under 8%.

It is also worth explaining the hill shapes of the improvement curves. At low arrival rates the DNN system times converge to those of NN. With large times between arrivals and many idle vehicles there is likely an idle vehicle close to a new customer and it is extremely unlikely for a vehicle assignment reroute to occur. As a result there is practically no difference in the behaviour of the two policies. But as the arrival rate increases the number of vehicle assignment reroutes also increases and at some point a difference between the policies can be seen. It should be noted that the greater the number of service vehicles the higher the critical traffic intensity point where the difference is first noticed. The greater the arrival rate the more assignment reroutes that occur in the DNN policy and the average s_w decreases as vehicles are picking up more newly arriving nearer customers. At some point the improvement reaches a maximum and then begins to decline because of the impact of a second factor. With larger arrival rates there are many customers awaiting service. It is likely that any newly idle cab will be near a currently waiting customer decreasing the chances of a beneficial vehicle assignment reroute. As the traffic intensity continues to increase fewer and fewer reroutes occur and the improvement declines.

4. ANALYTICAL APPROXIMATION OF POLICY PERFORMANCE

In addition to the simulated results, it would be valuable to derive analytically the expected policy performance. This is first done in detail for the FCFS policy followed by approximations of the NN policy.

4.1. Analytical Derivations for FCFS System Time

It is desired to develop analytical expressions to estimate the system times for the FCFS policy. The discussion begins with the single vehicle unit velocity case, followed by a varying velocity service vehicle and finally multiple service vehicles.

4.1.1. FCFS Single Service Vehicle with Unit Velocity

The expected service and system times for the single vehicle FCFS case are derived analytically. The approach is analogous to that used in Bertsimas and Van Ryzin (Bertsimas and Van Ryzin 1991), but modified for the Dynamic DARP case.

Larson and Odoni (1981) define geometric probability as follows: Given two uniformly and independently distributed points Y_1 and Y_2 in a square of area A, then

$$E[Y_1 - Y_2] = c_1\sqrt{A} \quad E[Y_1 - Y_2] = c_1\sqrt{A}$$

$$E[(Y_1 - Y_2)^2] = c_2 A \tag{2}$$

where $c_1 \sim 0.52$ and $c_2 = 1/3$

The expected service time can be written as follows

$$E[s] = E[s_w] + E[s_t] \tag{3}$$

where both $E[s_w]$ and $E[s_t]$ are equal to the expected distance between two uniformly distributed points in a square area and therefore follow Eq. (2) such that

$$E[s_w] = c_1\sqrt{A} = 0.52\sqrt{1} = 0.52$$
$$E[s_t] = c_1\sqrt{A} = 0.52\sqrt{1} = 0.52 \tag{4}$$

Then the expected service time is shown as follows

$$E[s] = c_1\sqrt{A} + c_1\sqrt{A} = 2c_1\sqrt{A} = 2(0.52)\sqrt{1} = 1.04 \tag{5}$$

Further, the analytical expression for the expected system time in the single vehicle FCFS case can be derived. Since

$$E[T] = E[W] + E[s] \tag{6}$$

and $E[s]$ is already known, to determine $E[T]$ the expression for $E[W]$ must be derived.

The single vehicle FCFS can be modeled as an $M/G/1$ queue (Bertsimas and Van Ryzin 1991). As a result, the well-known Pollaczek-Khinchin formula (Kleinrock 1976) can be used

$$W = \frac{\lambda s^2}{2(1-\rho)} \tag{7}$$
where s^2 is the second moment of the service time.

Using Eq. (2) the variance of the service time can be expressed as

$$Var[s] = Var[s_w] + Var[s_t] = 2Var[s_t] = 2(E[s_t^2] - E[s_t]^2) = 2(c_2 A - c_1^2 A) \tag{8}$$

From Eq.s (5) and (8) the second moment of the service time can be shown as

$$E[s^2] = 2c_2 A + 2c_1^2 A \tag{9}$$
$$E[s^2] = Var[s] + E[s]^2 = 2(c_2 A - c_1^2 A) + \left(2c_1\sqrt{A}\right)^2 = 2c_2 A + 2c_1^2 A$$

By inserting Eq.s (1) and (9) into (7) the expected waiting time is then

$$E[W] = \frac{\lambda(2c_2 A + 2c_1^2 A)}{2(1-2\lambda c_1\sqrt{A})} \tag{10}$$

Therefore, from Eq.s (5), and (10) it follows that the expected system time is

$$E[T] = \frac{\lambda(2c_2 A + 2c_1^2 A)}{2(1-2\lambda c_1\sqrt{A})} + 2c_1\sqrt{A} \tag{11}$$

4.1.2. FCFS Single Service Vehicle with Varying Velocity

The derivation is then extended to the cases where the service vehicle can travel with a constant velocity of magnitude *v*. The service time for the FCFS policy with a service vehicle travelling at velocity *v* is expressed as

$$E[s] = \frac{E[s_w] + E[s_t]}{v} \tag{12}$$

Then, from Eq. (4) it is easily shown that

$$E[s] = \frac{2c_1\sqrt{A}}{v} \tag{13}$$

Using an analogous approach as to that of Eq. (9) it is easily shown the second moment of the service time equals

$$E[s^2] = Var[s] + E[s]^2 = \frac{2(c_2 A - c_1^2 A)}{v^2} + \frac{\left(2c_1\sqrt{A}\right)^2}{v^2} = \frac{2c_2 A + 2c_1^2 A}{v^2} \tag{14}$$

Again, since FCFS can be modelled with an *M*/*G*/1 queue the waiting time can be calculated using Eq. (7) and thus the expected system time for the varying velocity FCFS policy can be expressed as follows

$$E[T_{FCFS-VV}] = E[W] + E[s] = \frac{\lambda E[s^2]}{2(1-\rho)} + \frac{2c_1\sqrt{A}}{v} = \frac{\lambda(2c_2 A + 2c_1^2 A)}{2(1-\rho)v^2} + \frac{2c_1\sqrt{A}}{v} \tag{15}$$

The analytical formula underestimates the simulated results by a small margin. But the relative error is never more than 5% which is close enough to consider the results to be in agreement.

4.1.3. FCFS Multiple Service Vehicles

The varying velocity formula can now be extended to estimate the system times for the multiple vehicle FCFS case. First, it is assumed that the ratio of the expected waiting time to the probability the waiting time is greater than zero for the multiple vehicle case is equal to the equivalent ratio for the varying velocity case. This is expressed as follows

$$\frac{E[W_{FCFS-Multi}]}{P\{W_{FCFS-Multi} > 0\}} \cong \frac{E[W_{FCFS-VV}]}{P\{W_{FCFS-VV} > 0\}} \tag{16}$$

The individual terms of Eq. (16) are each explained, beginning with the probability of the waiting time being greater than zero for the varying velocity case. This term is simply equal to the utilization and is therefore solved as follows

$$P\{W_{FCFS-VV} > 0\} = \rho = \lambda E[s] = \frac{2\lambda c_1\sqrt{A}}{v} \tag{17}$$

Furthermore, the waiting time for the varying velocity case is defined previously within Eq. (11). Lastly, the term representing the probability that the wait time is greater than zero for the multiple vehicle FCFS case is more difficult to solve. An exact solution is not known but it can be represented using *M*/*M*/*N* queue with the same λ and µ. The probability then becomes an Erlang C distribution expressed as

$$P\{W_{FCFS-Multi} > 0\} \cong C(N, \lambda) \tag{18}$$
where *N* is the number of service vehicles.

The expected waiting time for the multiple vehicle FCFS are then estimated as

$$E[W_{FCFS-Multi}] \cong C(N, \lambda) \times \frac{\frac{\lambda(2c_2 A + 2c_1^2 A)}{2(1-\rho)v^2}}{\rho} \tag{19}$$

And then the expected system time can be expressed as

$$E[T_{FCFS-Multi}] \cong C(N, \lambda) \times \frac{\frac{\lambda(2c_2 A + 2c_1^2 A)}{2(1-\rho)v^2}}{\rho} + 1.04 \tag{20}$$

For most cases the error between the simulated and analytical results hovers around 2% and even in the extreme case the maximum error is no more than 6%. This is certainly close enough to consider Eq. (20) to be an accurate analytical approximation of the expected system times for the FCFS policy with multiple service vehicles travelling at a constant unit velocity (see

Jagerman and Melamed, 2003 for the basis of this approximation).

4.2. Analytical Approximations for NN Single Service Vehicle

It would be desirable to also derive analytical expressions for the system times of the NN policy, but due to the unknown service distance this becomes a very difficult exercise. However, it is feasible to develop approximate analytical expressions based on coefficients derived from initial simulation runs. This is first done for the NN single vehicle varying velocity policy. The discussion begins with a derivation based on the unit velocity case and then the results are extended to the situation with a service vehicle travelling at a non-unit velocity.

4.2.1. NN Single Vehicle with Unit Velocity

In their work on the multiple m-vehicle infinite capacity DTRP, Bertsimas and Van Ryzin derive upper bounds for system time performance (Bertsimas and Van Ryzin 1993). Both light traffic and heavy traffic bounds are presented, with the heavy traffic bound being of interest here.

Bertsimas and Van Ryzin define the utilization of the service vehicle differently than in this paper. To avoid confusion the Bertsimas and Van Ryzin utilization is dubbed *r* and is defined as

$$r = \lambda E[s_t] \cong 0.52\lambda \tag{21}$$

The heavy traffic ($r \rightarrow 1$) lower bound can then be expressed as

$$T^* \geq \gamma^2 \frac{\lambda A}{m^2 v^2 (1-r)^2} - \frac{E[s_t](1-2r)}{2r} \tag{22}$$

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

While the γ constant derived for the DTRP cannot be applied directly, Eq. (22) can be suited to the Dynamic DARP.

First, for convenience, part of Eq. (22) is defined separately and is termed *K*,

$$K = \frac{\lambda A}{(1-r)^2} \tag{23}$$

It follows that there is a linear relationship between *K* and the system times for the single vehicle unit velocity NN policy. The R^2 value of 0.9986 validates the linear regression fit is a good one. Given this relationship it stands that with the correct coefficients for slope, *a*, and *y*-intercept, *b*, the system time can be accurately approximated for any λ using a linear equation. Motivated by Eq. (22), the *y*-intercept is fixed as the service time and the slope is determined from simulation results such that

$$E[T] = 0.515087K + 1.04 \tag{24}$$

The equation is verified against the simulation results. A maximum error of only 4% indicates the analytical approximation is valid.

4.2.2. NN Single Vehicle with Varying Velocity

The formula obtained from the unit velocity case can now be extended to the situation in which the service vehicle travels with a constant velocity of magnitude *v*. Simulated results could again be used to obtain the coefficients *a* and *b* for each individual velocity. However, the exercise becomes more useful if the same coefficient values apply independent of the velocity such that

$$E[T] = 0.515087\frac{K}{v^2} + \frac{1.04}{v} \tag{25}$$

This hypothesis is tested against the simulation results for the NN single vehicle policy with velocity increases of 10 and 20 times the unit velocity. Throughout all arrival rates the error stays less than 5% indicating that the analytical approximation provides an accurate estimate of the system time. Therefore, the hypothesis is indeed correct and the coefficients obtained from the unit velocity case can be used to calculate the expected system time for the NN policy with a single service vehicle travelling at a constant velocity of magnitude *v*.

4.3. Analytical Approximations for NN Multiple Service Vehicles

It is also desired to find analytical approximations for the NN policy with multiple service vehicles travelling at a constant unit velocity. An analogous approach to the varying velocity case is used.

Ideally, as in varying velocity, it would be best if the same coefficients could be derived from the single vehicle case and applied to all other vehicle cases. Unfortunately, from experimentation it is known that this does not hold true.

Although each number of vehicles requires its own set of coefficients, it can still be shown that analytical approximations can be derived for the NN multiple vehicle policy. As an example, the case of 10 service vehicles is presented. First, it is important to understand that the relationship between *K* and the system time is no longer linear. Therefore, a quadratic polynomial must be used for the regression fit. An R^2 value of 0.9916 indicates that the second order polynomial is indeed a good fit. Therefore, the equation used for the analytical multiple vehicle system time approximation is

$$E[T] = c\frac{K^2}{N^3} + d\frac{K}{N^2} + \frac{e}{N} \tag{26}$$

From simulation the coefficient values can be determined. For the *N=10* case the values are:

$$c = -0.000752, d = 0.67723, e = 7.3$$

Using the same method the coefficient values for other *N* cases can also be obtained. Here, the cases of $N=10$ and $N=20$ are examined. Again, the maximum relative error is less than 5%, further proving the regression fit is strong.

5. APPLICATION OF POLICIES IN CITY-LIKE ENVIRONMENT

The relative performance of the policies is studied under somewhat more realistic conditions, which is dubbed the "City" environment. Although this "City" model is a large simplification of an actual city, it is still believed some insight can be gained into how the policies may perform in the real world.

5.1. Description of "City" Environment

The unit square service region is divided into two halves along the vertical midpoint line. The left half is termed the *city* and the right half the *suburbs*. Both the *city* and *suburb* sections have areas of $A_c=A_s=0.5$. The total area of the region is still $A=1$.

The simulation time is divided into repeating 24 hour time chunks representing a fictitious *day*. The first 12 hours of the *day* are termed the *morning* and the second 12 hours the *evening*.

The "City" environment is constructed to emulate a simplified form of a typical city's traffic patterns. During the *morning* 50% of the demands have a pick-up location in the *suburbs* and a drop-off location in the *city*. The pick-up locations are uniformly distributed within the *suburbs* and the drop-off locations are independent of the pick-up locations and uniformly distributed within the *city* area. The other 50% of demands have pick-up and drop-off locations independent of the *city/suburb* divide as in the base case. Furthermore, there is an emulated *rush hour* period between hours 7 and 10 where the arrival rate of demands is doubled. In the *evening* the pattern is reversed such that 50% of the demands have pick-up locations in the *city* and drop-off locations in the *suburbs* and the other 50% remain independent of the *city/suburb* boundaries. The *evening rush hour* with the doubled traffic rate occurs between hours 16 and 19.

5.2. Comparison between "City" Policies

The results are compared between the policies under the "City" conditions.

5.2.1. Comparison of NN "City" to FCFS "City"

The relative performance of the NN and FCFS policies in the city-like conditions is summarized in Figure 4.

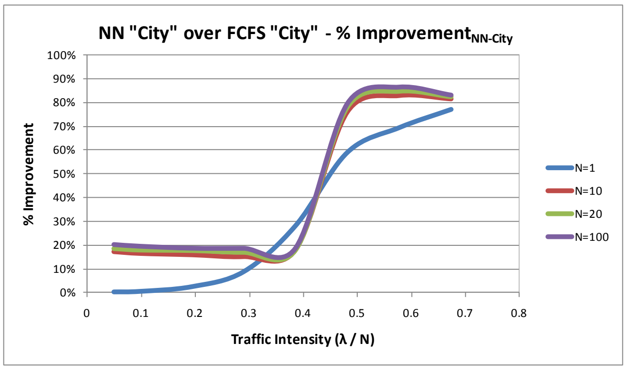

Figure 4: System Time Percent Improvement for NN "City" Over FCFS "City"

As in the base case the NN policy outperforms FCFS in the "City" environment. The improvement follows a similar shape but the differences between the number of vehicles is less pronounced. In fact, there is virtually no difference between any of the multiple vehicle cases. Furthermore, the magnitude of the improvement is slightly less under the "City" conditions. At low arrival rates the improvement hovers around 20% compared to values of 25-40% for the base case and the asymptotic behaviour is less severe in the "City" environment.

5.2.2. Comparison of DNN "City" to NN "City"

Similarly, the relative performance of the DNN and NN "City" policies is shown in Figure 5.

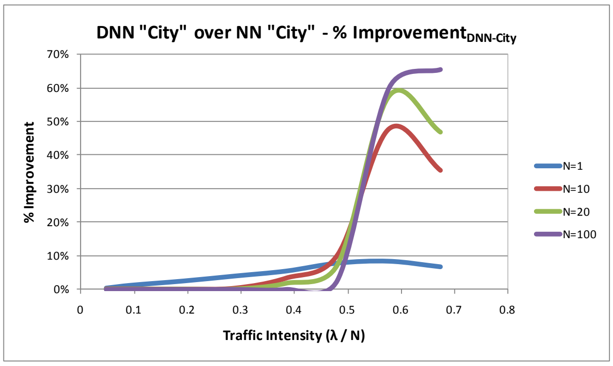

Figure 5: System Time Percent Improvement for DNN "City" Over NN "City"

The DNN "City" policy outperforms NN "City". However, here the pattern is quite a bit different from the base case. For lower traffic intensities the improvement is non-existent to very small, but at higher arrival rates the improvement spikes to 40-60% dwarfing the maximum improvement of only 10% in the base case.

The primary conclusion from the analysis is that the relative order of the policies remains unchanged in the city-like environment.

6. ANTICIPATORY BEHVAVIOUR IN CITY-LIKE ENVIRONMENT

The dynamic routing policies studied thus far show dependable performance improvements, even in the city-like conditions. However, these policies are all reactive in nature in that routing decisions are not made until a new customer arrives. What if, instead of idle taxis remaining in their current location waiting for the next call, idle taxis are routed to locations that are most likely to see new customers? Would this result in further performance improvements? Such a policy is what is referred to as anticipatory vehicle routing or, more generally, anticipatory behaviour. In essence, the goal of anticipatory behaviour is to predict, with some uncertainty, how to best distribute the vehicles throughout the service area to most efficiently serve the upcoming demands.

6.1. Description of Anticipatory Behaviour Model

The details of the anticipatory model studied are now described. The algorithm is designed in an attempt to optimize the number of service vehicles present in each sector to best meet the upcoming demand profile. For example, in the *morning* the flow of demands is from the *suburbs* to the *city*. This leaves an excess of vehicles in the *city*. Therefore, on completed service, a portion of the service vehicles are directly rerouted to the *suburbs*. In order to optimize the anticipatory routing, it is crucial that the correct proportion of vehicles be rerouted. Such a derivation of this proportion is now presented.

Consider the *morning* commute. Recall that 50% of the demands have pick-up locations in the *suburbs* and drop-off locations in the *city*, and the other 50% have random pick-up and drop-off locations. Since half of the random drop-off locations will be in the *city*, overall 75% of the vehicles will end in the *city* and only 25% in the *suburbs*. Conversely, 75% of the demands will have pick-up locations in the *suburbs* and only 25% in the *city*. With no anticipatory routing the mismatching of service vehicle locations to the demand profile is apparent.

To properly match the demand profile 75% of the total number of vehicles should be present in the *suburbs*. With 25% of the taxis already ending in the *suburbs*, 50% of the total number of service vehicles should be routed from the *city* to the *suburbs*. Therefore, in the *morning*, of the service vehicles with a drop-off location in the *city* 66.67% are rerouted back to the *suburbs* even if no specific demand is waiting. It is randomly determined which of the vehicles are rerouted. In the *evening* the logic remains the same but the pattern is reversed.

It is important to note that the preceding rules only apply when there are no customers waiting. If a service vehicle becomes newly idle and there are one or more waiting demands the vehicle is immediately sent to service a waiting demand, and therefore the anticipatory routing no longer applies. Furthermore, idle service vehicles which are travelling towards a new anticipatory location are still eligible to be assigned to incoming demands. The intermediate position of a travelling idle vehicle is dynamically updated and compared to the location of the incoming demand as would be done if the service vehicle were not moving.

An additional point centers around the policies studied with anticipatory vehicle routing. It does not seem logical to apply FCFS with anticipatory behaviour because the service order is based on arrival time rather than relative location. As a result the policy is not studied here. The NN and DNN policies are considered. Both policies behave exactly as before except for the fact that service vehicles are anticipatorily routed to better meet the upcoming demand patterns.

To route the idle vehicles an additional feature is added to the "City" conditions. Cab-stops are introduced such that when a service vehicle becomes idle it does not remain at its current position but travels to a cab-stop. The cab-stops emulate one additional real-life feature in that idle vehicles do not normally wait where their last drop-off was but instead at designated taxi waiting areas. In the model there are a total of eight cab-stops – four in the *city* and four in the *suburbs* – aligned vertically through the center of each sector. The cab-stop locations are chosen such that a uniformly distributed newly idle service vehicle has an equal probability of being nearest each of the eight stops. Each cab-stop encompasses a nearest vehicle area with a width of 0.5 and a height of 0.25.

Given the overall profile of anticipatory vehicle routing, the problem then becomes to which particular cab-stop to send each service vehicle. If a vehicle is not being rerouted to the opposite sector, it will simply travel to the nearest cab-stop. If the vehicle is being anticipatorily routed it will travel to the nearest cab-stop in the opposite sector. For example, in the *evening*, a vehicle located in the *suburbs* at (0.82, 0.46) would be sent to the nearest *city* cab-stop at (0.25, 0.375).

6.2. Comparison of Anticipatory Policies to "City" Policies

The primary goal of studying anticipatory routing is to understand the improvement, if any, it generates over the reactive "City" policies. This comparison is shown for the NN and DNN policies in Figures 6 and 7, respectively.

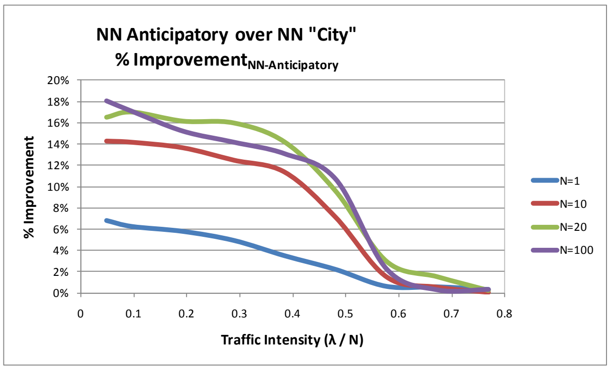

Figure 6: System Time Percent Improvement for NN Anticipatory Over NN "City"

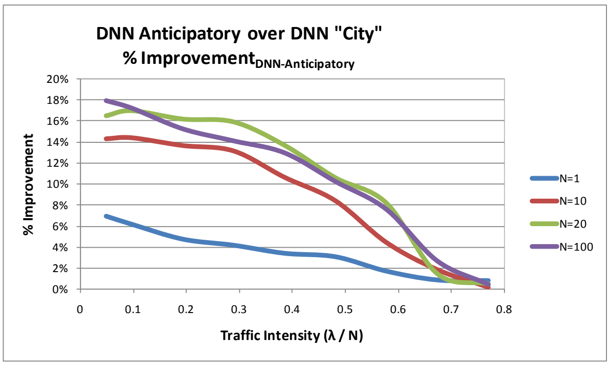

Figure 7: System Time Percent Improvement for DNN Anticipatory Over DNN "City"

The first note is that the improvement for the NN and DNN policies is very similar. This provides evidence that the effect of adding anticipatory behaviour does not significantly depend on to which of the two policies it is applied. Therefore, the performance of anticipatory behaviour is discussed in general terms.

Anticipatory vehicle routing does outperform the reactive "City" policies. While the magnitude of the improvement does not depend on the policy, it is affected by the traffic intensity and the number of service vehicles. The maximum improvement of around 18% occurs at the lowest traffic intensities. The magnitude of the improvement remains relatively constant up to medium traffic before it begins to decline more quickly. In heavy traffic the system times are statistically equal to those of the "City" policies. Furthermore, the improvement is greater for the multiple vehicle cases than the single vehicle case. The maximum improvement for $N=1$ is only around 7%. The 20 and 100 vehicles cases outperform the 10 vehicle case but the improvements for $N=20$ and $N=100$ are quite close. This suggests the magnitude of the improvement increases with number of service vehicles up to a point where it then begins to level off.

These results are not entirely surprising. It seems logical the improvement would decrease as the arrival rate increases. At low traffic intensities rerouted idle service vehicles have sufficient time to reach or make significant progress towards the opposite sector before they are assigned to an incoming demand. As the arrival rate increases idle taxis are assigned to new demands more quickly and on average have less travelling time towards the other sector, thus lowering the benefit. At high traffic intensities there is almost always a waiting demand and thus the anticipatory policy collapses to the reactive "City" policy and no improvement is seen.

7. CONCLUSION

This paper makes several contributions to the understanding of service vehicle routing in the Dynamic DARP. First, the well understood NN policy was extended and modestly improved to incorporate partially dynamic behaviour. It was shown that this DNN policy outperformed NN by up to 10% under the base conditions. Secondly, the analytical understanding of the policies was furthered. Analytical formulas for the multiple vehicle FCFS policy were derived. Furthermore, accurate approximations for the single vehicle varying velocity NN policy were presented using only a single set of simulation runs. While not yet completed, this potentially paves the way for similar expressions for multiple vehicle policies. Thirdly, it was demonstrated that the relative performance of the FCFS, NN, and DNN policies holds true under somewhat more realistic city-like conditions. Lastly, the understanding of anticipatory behaviour for service vehicle dispatching was furthered. It was shown that anticipatory vehicle routing for the more complex NN and DNN policies outperforms the reactive policies by up to 18%.

8. FUTURE WORK

The major areas of potential future work are discussed next.

8.1. Future Work for Comparison of Base Case Policies

The DNN policy is not a fully dynamic solution. Two of the major shortcomings include: (i) if a new demand arrives and there is one or more idle vehicles, the demand is always serviced by an idle vehicle even if an assigned vehicle is nearer, and (ii) no dynamic reassignments occur when a service vehicle becomes newly idle after reaching a drop-off location. There is a significant opportunity to develop and test a fully dynamic nearest neighbour policy. It is hoped that such a policy could produce much greater improvements than the current magnitudes of no more than 10%.

While there are many potential solutions to the implementation of a fully dynamic policy, one possibility is discussed here. The policy is concisely summarized as: When a service vehicle (demand) becomes idle (arrives) it is assigned to its nearest demand (service vehicle), and dropped service vehicles (demands) are considered as becoming idle (arriving). Whenever a service vehicle becomes idle or a new service vehicle arrives, the service vehicle-demand pairs are dynamically reassigned. This is done in such a way that once a vehicle and demand have been paired they are removed from future comparisons for that trigger such that algorithm is guaranteed to quickly converge to a solution

8.2. Future Work for Analytical Approximations of Policy Performance

As the policies become more dynamic and the service environments more complex, producing analytical expressions for system time performance becomes increasingly more difficult. With that being said, it is still believed there is an opportunity to develop an analytical approximation for the multiple vehicle NN policy that requires only one set of initial simulation runs. It is unclear exactly how this is to be accomplished but the single vehicle varying velocity NN model could be used as a starting point.

8.3. Future Work for Objective 4: Application of Policies in "City" Environment

The "City" environment studied was a very simple approximation of the real-world. The conditions could be expanded to better simulate real life conditions. Potential improvements could include: the addition of roads to constrain taxi movement, an increased number of neighbourhoods, a larger number of service vehicles, more complex traffic patterns and rush hour modelling, or traffic disturbances such as stoplights. Although many improvements can be made it seems unlikely that any model could accurately simulate the complexity of a real city so the models should still only be used to assess relative, not absolute, performance. That being said, there is still value in understanding the effects that different real-life phenomenon could have on policy performance.

8.4. Future Work for Anticipatory Behaviour in "City" Environment

Anticipatory vehicle routing represents a significant opportunity for future work. There are many possible directions to take in future research, so the ones presented here are merely suggestions. The effect of the vehicle anticipation routing method could be further studied. For example, instead of travelling to a particular location, taxis could be sent to patrol throughout the specified sector. Furthermore, anticipatory behaviour could be applied to the previous future work suggestions. It would be interesting to understand how the performance of anticipatory routing would change under more realistic "City" conditions. Furthermore, anticipatory behaviour could be applied to a fully dynamic policy. If possible, it would be of interest to understand how close such a policy would come to the optimal system performance.

REFERENCES

Berbeglia, G., Cordeau, J.F. and Laporte, G., 2010. Dynamic pick-up and delivery problems. *European Journal of Operations Research*, 202 (1) 8-15.

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

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

Jagerman, D.L. and Melamed, B. 2003. Models and Approximations for Call Center Design. *Methodology and Computing in Applied Probability*, 5, 159-181.

Kleinrock, L., 1976. *Queuing Systems, Vol. 1: Theory*. New York: Wiley.

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

Ozesenli, O.N., and T. Demirel. 2005. A dynamic approach (SAT) to vehicle allocation problem with uncertain demands for service sector. *35th International Conference on Computers and Industrial Engineering. Istanbul, Turkey.*

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

Xiang, C., C. Chu, and H. Chen. 2008. The study of a dynamic dial-a-ride problem under time-dependent and stochastic environments. *European Journal of Operations Research*, 185 (2) 534-551.

Xu, H., 1994. *Optimal policies for dynamic and stochastic vehicle routing problems*. Thesis (PhD). Massachusetts Institute of Technology.

AUTHORS BIOGRAPHY

Thomas Pickles is a 4th year mechanical engineering undergraduate student in the Department of Mechanical and Industrial Engineering at the University of Toronto. His research interests are in the areas of discrete event simulation and maintenance optimization. He has won several scholarships for academic excellence including University of Toronto Scholar, Professional Engineers of Ontario Scholar, and Natural Sciences and Engineering Research Council of Canada scholarship.

Barış Balcıoğlu is an Associate Professor in the Department of Mechanical and Industrial Engineering at the University of Toronto. He received his Ph.D. in Industrial and Systems Engineering from Rutgers University in 2003. His research and teaching activities are in the areas of queueing theory, discrete-event simulation, performance analysis of call centers and production/inventory systems, and reliability and maintenance engineering. His research has been supported by the Natural Sciences and Engineering Research Council of Canada.

Beno Benhabib received his Ph.D. degree in mechanical engineering from the University of Toronto, Canada, in 1985. Since 1986, he has been a Professor of mechanical and industrial engineering at the University of Toronto, with cross appointments to the Department of Electrical and Computer Engineering as well as to the Institute of Biomaterials and Biomedical Engineering. His current research interests include development of autonomous robotic systems for manufacturing, biomedical, and service industries.