

Waiting Strategies for Dynamic Vehicle Routing

Jürgen Branke

Institute AIFB, University of Karlsruhe, 76128 Karlsruhe, Germany, branke@aifb.uni-karlsruhe.de

Martin Middendorf

Department of Computer Science, University of Leipzig, 04109 Leipzig, Germany, middendorf@informatik.uni-leipzig.de

Guntram Noeth

McKinsey & Company, Inc., Prinzregentenstr. 22, 80538 Munich, Germany, guntram_noeth@mckinsey.com

Maged Dessouky

Department of Industrial and Systems Engineering, University of Southern California, Los Angeles, California 90089, maged@rcf.usc.edu

Many real-world vehicle routing problems are dynamic optimization problems, with customer requests arriving over time, requiring a repeated reoptimization. In this paper, we consider a dynamic vehicle routing problem where one additional customer arrives at a beforehand unknown location when the vehicles are already under way. Our objective is to maximize the probability that the additional customer can be integrated into one of the otherwise fixed tours without violating time constraints. This is achieved by letting the vehicles wait at suitable locations during their tours, thus influencing the position of the vehicles at the time when the new customer arrives. For the cases of one and two vehicles, we derive theoretical results about the best waiting strategies. The general problem is shown to be NP-complete. Several deterministic waiting strategies and an evolutionary algorithm to optimize the waiting strategy are proposed and compared empirically. It is demonstrated that a proper waiting strategy can significantly increase the probability of being able to service the additional customer, at the same time reducing the average detour to serve that customer.

Key words: dynamic vehicle routing; waiting strategies; flexible operation

History: Received: July 2002; revisions received: August 2003, March 2004; accepted: March 2004.

1. Introduction

Dynamic vehicle routing with previously unknown customer requests arriving over time has become increasingly important in the transportation industry as new technologies such as global positioning systems and wireless communications allow the assignment of new requests to vehicles in real time. An optimization algorithm in such a dynamic environment has to repeatedly adapt a solution whenever new information becomes available. If most customer requests are known in advance, a simpler and more practical approach is to design preplanned routes for the known customers, and to insert the few new customers as they become known into the already planned routes. However, if new requests are expected, rather than just reacting to the new demand, one should anticipate a change by trying to maintain flexibility (Branke and Mattfeld 2000). As we will show, for the vehicle routing problem (VRP) such flexibility can be maintained by having the vehicles wait at appropriate locations in their tours.

In this paper, we consider a VRP with a single new customer arriving after the vehicles have left the depot. The location of the new customer is assumed

to be uniformly distributed within the service region, and we consider the cases of either a known or unknown arrival time. We assume that there is slack time in the preplanned tours that can be used to insert a new customer request or to wait, or both. The new customer may be inserted between two subsequent customers of any one of the given tours, as long as the vehicle has enough slack time. If no vehicle can integrate the new customer, the new customer is not serviced.

We examine the problem of finding an optimal waiting schedule for the vehicles to maximize the probability that a new customer can be incorporated into one of the tours. Intuitively, serving customers as quickly as possible (i.e., without waiting) seems to be a good strategy. Waiting has the clear disadvantage of using up precious time that cannot be used later to make a detour to serve the additional customer. Indeed, we prove that for the case of a single vehicle, not to wait is the optimal strategy. In the case of several vehicles, however, waiting may be beneficial, because it allows the vehicles to remain at strategically favorable locations. We prove that for the general problem with more than one vehicle the problem

of finding an optimal waiting schedule is NP-hard. We present and compare several heuristics and an evolutionary algorithm, and show empirically that a proper waiting strategy improves the probability of being able to service an additional customer, while reducing the average length of the detour that is necessary to integrate this customer.

We will also provide evidence that the benefits of waiting will prevail even when more than one new customer is to be inserted, as long as the number of new customers is small. When the number of new customers is large, waiting is unlikely to be of any benefit, but the strategy of working with preplanned tours becomes questionable anyway.

The remainder of the paper is organized as follows. In §2 we survey some related work. The definition of the considered dynamic vehicle routing problem and the NP-completeness results can be found in §3. In §4, we show theoretical results for the cases of one and two vehicles. Several heuristics for the general case are suggested in §5. The heuristics are analyzed empirically in §6. Some problem variants, in particular the case of more than one new customer, are discussed in §7. The paper concludes with a summary and an outlook for future work in §8.

2. Related Work

Dynamic VRP has been studied extensively in the literature, with different aspects of the problem changing over time. An introduction to dynamic vehicle routing, explaining the differences between static and dynamic vehicle routing, has been given by Psaraftis (1988, 1995). A recent classification and survey on dynamic vehicle routing has been provided by Bianchi (2000). Further surveys can be found in the works of Gendreau and Potvin (1998) and Bertsimas and Simchi-Levi (1996). Because our approach is concerned with new customers arriving over time, the remaining section is restricted to related work dealing with that particular aspect of dynamism.

In the literature, one can find two main approaches to cope with the problem changes caused by new customer arrivals. The two strategies are reoptimization and dispatching. The first approach is to reoptimize the VRP whenever a new customer arrives. A typical example is the tabu search algorithm for the dynamic VRP with time windows presented by Gendreau et al. (1996, 1999). In the proposed approach, a tabu search heuristic is running continually, trying to improve the current best solution. It is interrupted only by two events: Either a new request arrives that has to be scheduled, or a customer has to be removed because it is currently served by a vehicle. If a new request arrives, it is inserted into the existing solution at the location that minimizes a weighted sum of

detour and service delay, then a local search heuristic is performed until a local minimum is found, and the tabu search is resumed. Once a customer of a vehicle has been served, the best solution found up to now is used to determine the next stop for the particular vehicle and all other current solutions are updated correspondingly. Based on the same tabu search heuristic, in a subsequent paper, Ichoua, Gendreau, and Potvin (2000) examine the benefit of allowing a vehicle to change its destination while under way from one scheduled customer to the next in order to service a newly arrived customer. Empirical tests show a reduction in the number of unserved customers and in the combined objective function addressing the distance traveled and the total late-ness, if diversion is allowed.

Diversion is investigated also by Regan, Mahmassani, and Jaillet (1995c) in the context of a dynamic single-vehicle, single-capacity pickup and delivery problem. Using simulation, they conclude that diversion can reduce the overall distance traveled. However, if the system has to deal with multiple dynamically arriving requests, the number of diversions must be limited to avoid long travel distances toward the end of the service horizon. Regan, Mahmassani, and Jaillet (1995a) extend this work to a multiple vehicle environment, and Regan, Mahmassani, and Jaillet (1995b) investigate the effect of allowing reassignment of customers that are already scheduled but not yet served.

Yang, Jaillet, and Mahmassani (1999) compare an exact branch-and-cut algorithm solved every time the problem changes, with a resequencing and reassignment heuristic. As might be expected, the constructive algorithm suffers from lack of computation time if the congestion level is high, and therefore performs worse than the heuristic approach. However, even for low congestion levels the heuristic rescheduling approach gives almost as good results as the constructive algorithm. Similarly, Powell, Towns, and Marar (1998) present a model that allows the human dispatcher, if he so desires, to override a solution provided by the system. They show that in a dynamic environment greedy heuristics frequently produce more usable solutions than algorithms guaranteeing mathematical optimality.

The second way to handle problem dynamics is to renounce planning to a certain extent, and, instead of scheduling all known tasks, only decide on a vehicle's next task. In this case, whenever a customer has been serviced and a vehicle is available, dispatching rules are used to decide which customer should be assigned to the vehicle next. An example for the use of this approach is the work by Bertsimas and van Ryzin (1991, 1993) on the dynamic traveling repairman problem, where queuing models are used

to compare the impact of different dispatching rules on the average time that a customer spends in the system (while waiting for service or being served). For the uncapacitated one-vehicle case, Bertsimas and van Ryzin (1991) derive an optimal dispatching policy for light traffic and a lower bound for the expected waiting time in heavy traffic. Papastavrou (1996) proposes for this problem a policy that performs optimal in light traffic and does not perform worse than twice the lower bound in heavy traffic. Bertsimas and van Ryzin (1993) and Swihart and Papastavrou (1999) extend this approach of developing dispatching rules by applying queuing theory to a wider scope of problem areas. Bertsimas and van Ryzin (1993) study the dynamic traveling repairman problem with multiple capacitated and uncapacitated vehicles, and Swihart and Papastavrou (1999) apply this methodology to the single-vehicle pickup and delivery problem.

Instead of just reacting to problem changes (e.g., the arrival of new customers), however, it may be beneficial to anticipate such events by positioning the vehicles in strategic locations. Some first attempts in that direction can be found in the literature. In the aforementioned paper by Bertsimas and van Ryzin (1991), it is shown that in the case of a single vehicle and customers appearing uniformly distributed in an Euclidean plane, it is optimal to reposition the vehicle to the center of the service region whenever there are no customers left to be serviced. Returning to the center anticipates future customer arrivals by positioning the vehicle so that the expected distance to the next arriving customer is minimized. Of course, this strategy assumes light traffic, with the vehicle being mostly idle. Similar results are shown in Bertsimas and van Ryzin (1993) for the light traffic case of the multiple-vehicle traveling repairman problem and in Swihart and Papastavrou (1999) for the single-vehicle pickup and delivery problem. Kilby, Prosser, and Shaw (1998) propose for a VRP that a vehicle should not drive back to the depot when idle, but should wait at the last customer location. This, of course, increases the probability of being able to service new customers in the area of the last customer. Powell (1986, 1996) studies a dynamic assignment problem where a fleet of vehicles is assigned to a set of locations with dynamically occurring demands. He shows that it is advantageous to take forecasted demands into account when deciding where the vehicles should drive next, compared with a model that only reacts after new demands have arrived. This, however, assumes that the demands can be accurately predicted. Thomas and White (2004) study anticipatory routing in the case of a few known locations of potential customers that might issue a request while the vehicle is under way.

In this paper, we examine the question of where vehicles should wait to maximize the probability that a new customer, appearing anywhere in the service region, can be integrated into one of the tours. Other than the previously described approaches for vehicle routing, we do not restrict such an anticipatory action to the case when vehicles are idle, but allow a vehicle to wait, even if it has not yet serviced all known customers. As we will show, allowing the vehicles to wait can greatly improve the service probability for the new customer.

Recently, Mitrovic-Minic and Laporte (2004) have also looked at the benefit of waiting strategies. The authors examine whether waiting strategies can reduce the total detour or the number of required vehicles for a dynamic pickup and delivery environment with time windows. They show that if all the available time is used for waiting at the beginning, total detour can be reduced, but more vehicles are needed to serve the additional customers. If, on the contrary, all waiting time is used at the end, the total number of vehicles can be reduced at the cost of having larger detours. Therefore, they propose a mixture of both strategies where the total tour is partitioned into service zones and the total time available for waiting is spread among them proportional to the time necessary to serve each service zone. With this waiting heuristic, they are able to reduce the detour and the required number of vehicles at the same time. Mitrovic-Minic, Krishnamurti, and Laporte (2004) combine the idea of waiting with the issue that short-term objectives (the next two hours) and long-term objectives (the remaining part of the day) are different in dynamic routing environments; the authors apply a double horizon-based heuristic with different objectives for these time periods.

The above-mentioned work has been carried out independently from our paper. The main differences are as follows: First, we consider a standard VRP, whereas Mitrovic-Minic and Laporte (2004) consider a pickup and delivery problem with time windows. Second, our aim is to maximize the probability of being able to service an additional customer, but they focus on the expected detour. In addition to experimental results, we also provide some theoretical analyses. Nevertheless, it is interesting to note that their results regarding “wait at the beginning” and “wait at the end” are consistent with the computational results presented in §6 of this paper.

3. Dynamic Vehicle Routing and the Waiting Drivers Problem

3.1. Problem Definition

Given are customers $1, 2, \dots, n$ and a depot 0 (also called customer 0) in the Euclidean plane, a fleet of m

vehicles, and for every vehicle i , $i \in [1:m]$ a tour $r_i = (0, c_{i1}, c_{i2}, \dots, c_{in}, 0)$, $c_{ij} \in [1:n]$, $c_{ij} \neq c_{kl}$ for $ij \neq kl$, i.e., a sequence of customers that originates and terminates at the depot. Let d_{ij} , $0 \leq i < j \leq n$ be the Euclidean distance between customers i and j . All vehicles depart from the depot at time 0 and must be back at the depot at time $T > 0$ at the latest. It is further assumed that vehicles drive on a direct line from one customer to the next with the same constant speed. Without loss of generality, we assume that the speed is one distance unit per time unit so that we can identify distance with travel time. Furthermore, there are no service times at a customer's location. Then, the total length of a tour r equals the time t_r to finish the tour so that all customers have been serviced. Time $w_r = T - t_r \geq 0$ is called the slack of the tour and can be used to wait at customer locations or to make a detour to serve an additional customer. It is assumed that all vehicles service their customers in the order as defined by their tour, except for a possible insertion of a new customer between two existing customers. Note that we allow a vehicle to change its direction during its drive from one customer to the next when the new customer becomes known and such a diversion seems appropriate.

In this paper, we study the case that exactly one additional customer request arrives at a uniformly distributed random location within a given convex service region. If several vehicles could service the customer, the vehicle with the smallest required detour will service it. If insertion is not possible into any tour (because otherwise the tour would take longer than time T), the customer request is rejected and the customer is not serviced.

A *waiting strategy* for a tour r is an assignment of waiting times to the customers of r (including the depot) such that the sum of the waiting times is less than or equal to the tour's slack w_r . When waiting time δ is assigned to the depot this means the vehicle starts at the depot at time $\delta > 0$. A waiting strategy for a set of tours R consists of waiting strategies for all tours r in R . The strategy to not assign any waiting times to the customers is called *NoWait*. A location x in the plane is said to be t_c -covered by a waiting strategy, if and only if the waiting strategy allows the integration of a new customer at location x arriving at time t_c into one of the given tours. Otherwise the location is called t_c -uncovered. An area is t_c -covered (t_c -uncovered) when every location in the area is t_c -covered (t_c -uncovered). Location x is called *not reachable* when there is no waiting strategy for R such that x is covered. A region in the plane is *not reachable* when every location in the region is not reachable. Note that at any time t_c , for a set of tours R and a waiting strategy for R , the probability to serve a customer equals the ratio of the t_c -covered area of the

service region to the size of the whole service region. We call the service region *restrictive* when it does not include all reachable locations. Otherwise it is called *nonrestrictive*.

Assuming that the additional customer may appear at any time $t_c \in [0..T]$ within the service region, we are looking for an optimal waiting strategy, i.e., a waiting strategy that maximizes the probability of servicing the new customer. Formally, we define the corresponding decision problem for the case that t_c is known in advance.

Waiting drivers problem (WDP) with known customer arrival time:

Given: Customers $0, 1, 2, \dots, n$ in the Euclidean plane where customer 0 is the depot, and a convex service region. A set R of m tours so that all customers are serviced by these tours. Positive constants $p_0 \leq 1$ and $t_c \in [0, T]$, where $T > 0$ is the total time available.

Question: Is there a waiting strategy for R so that the probability to serve the new customer arriving with uniform random distribution within the service region at time t_c is at least p_0 ?

The corresponding problem where the arrival time t_c of the customer is a random variable that is uniformly distributed in $[0, T]$ and is not known in advance is called the *WDP with unknown customer arrival time*.

3.2. NP-Hardness

In this section, we show the NP-hardness of WDP even under severe restrictions that are relevant for practical applications. The first theorem shows that WDP is NP-complete, no matter whether the customer arrival time is known or unknown, and even when the service region is nonrestrictive. Hence, no polynomial time algorithm exists for the WDP (if $P \neq NP$). The proof for the first theorem uses a reduction of the 3-SAT problem and is given in the appendix.

THEOREM 1. *The WDP with known or unknown customer arrival time is NP-complete even when the service region is nonrestrictive.*

For many practical applications, each individual tour contains only few customers and the total service region is divided into subregions so that each subregion is serviced by its own tour. Hence, the question arises whether the WDP problem becomes polynomial time solvable under such additional restrictions. Unfortunately, we can show with Theorem 2 that WDP remains NP-complete even when all tours are simple in the following sense. Each tour contains only a constant number of customers and the main parts of the tours between the first customer and the last customer do not cross each other. Only when a vehicle is on the way from the depot to the first customer or back to the depot might it cross other tours.

Because the proof of Theorem 2 is rather technical, it is omitted here. The proof’s idea is to use a suitable NP-complete planar version of the 3-SAT problem where each variable occurs at most three times positive and at most three times negated in a clause, and where the subgraphs that represent the variables are bipolar (Lichtenstein 1982).

THEOREM 2. *The WDP with known or unknown customer arrival time is NP-complete even when the service region is nonrestrictive and every tour contains at most a constant number of customers, and where all subtours between the first and the last customer do not cross each other (i.e., for each tour it is allowed that the part between the depot and the first customer and the part between the last customer and the depot crosses other tours).*

4. Optimal Waiting Strategies for Special Cases

In this section, we derive optimal waiting strategies for some special cases of the WDP. We prove that for a single vehicle and when the service region is nonrestrictive the *NoWait* strategy is optimal. When one of these conditions does not hold (i.e., the service region is restrictive or there is more than one vehicle), waiting can be advantageous even for simple problems.

4.1. Single Vehicle

We start with an example that shows that waiting can be advantageous for WDP with a single vehicle when the service region is restrictive, even when all customers lie within the service region. In the example, we assume that the arrival time is known. Similar reasoning shows that the result holds also for an unknown arrival time. Let $r = (0, 1, 0)$ be the tour of the vehicle where customer 1 has distance 1 from the depot. Let $T = 4$ and $t_c = 2$. First, consider the case when the vehicle does not wait. It is back at the depot at time $t_c = 2$. Clearly, then the area that is t_c -covered by the nonwaiting vehicle is a circle of Radius 1 around the depot (see Figure 1). For the

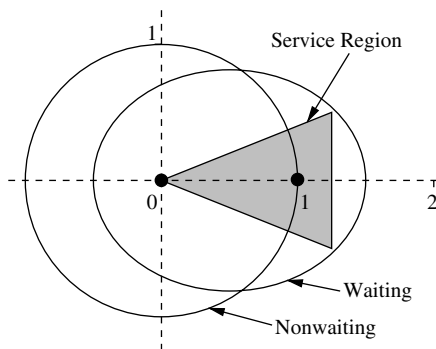


Figure 1 Example for a Single Vehicle: Service Region and Areas t_c -Covered by the Waiting and the Nonwaiting Vehicle

case of a waiting vehicle, assume it waits one time step at customer 1. Then, the vehicle is at the location of customer 1 at time $t_c = 2$. It can service an additional customer at any location for which the sum of the distances to customer 1 and the depot is at most two. Assuming that the locations of the depot and customer 1 have coordinates $(0, 0)$ and $(1, 0)$, respectively, the t_c -covered area is an ellipse of the form $x^2 + (4/3)y^2 = 1$. Note that the depot and customer 1 lie on the focal points of the ellipse, its major semi-axis has length 1, and its minor semi-axis has length $\sqrt{3}/2$. Assume that the service region is a triangle with vertices $(0, 0)$, $(5/4, 1/2)$, $(5/4, -1/2)$. It is easy to see that the whole triangle is t_c -covered by the waiting vehicle but not by the nonwaiting vehicle. The following theorem shows that waiting is never advantageous for WDP with a single vehicle when the service region is nonrestrictive.

THEOREM 3. *For the WDP with known and unknown customer arrival time, a single vehicle and a nonrestrictive service region, the optimal waiting strategy is NoWait.*

PROOF. First, assume that the new customer arrives at known time t_c . Let $T > t_c$ be the total service time. For a contradiction, let us assume that it is optimal to wait at some location (locations) before time t_c and consider a waiting vehicle with an optimal waiting strategy and a nonwaiting vehicle. Let $\delta > 0$ be the total waiting time of the waiting vehicle before t_c . Without loss of generality, we assume that $(x, 0)$ and $(y, 0)$, $y \leq x$ are the positions at time t_c of the nonwaiting vehicle and the waiting vehicle, respectively.

Let us at first assume that there is no customer at location $(x, 0)$. Let $1, 2, \dots, n$ be the sequence of customers that remain to be served by the nonwaiting vehicle after time t_c in this order, and let z_1, \dots, z_n be their locations. The depot is the final customer 0 with location z_{n+1} . Without loss of generality, we assume that the next customer to be serviced by the waiting vehicle and the nonwaiting vehicle is the same, and thus the sequence of customers that have to be serviced by the waiting vehicle after time t_c is the same as for the nonwaiting vehicle. Otherwise, there would be an additional customer on the tour of the waiting vehicle before it would reach location $(x, 0)$, which would reduce the covered area of the waiting vehicle, but does not influence the covered area for the nonwaiting vehicle. Then, $x - y = \delta$ holds and the location of the first customer on the tour after $(x, 0)$ is of the form $z_1 = (x', 0)$ with $x' > x$. Let $d := x' - x$ be the distance between $(x, 0)$ and $(x', 0)$.

Because the new customer can only be integrated into the tour without changing the sequence of the other customers, it follows that the total covered area is the union of areas that are covered by a set of ellipses that are defined in the following. Let

E_1, \dots, E_{n+1} be the ellipses corresponding to the non-waiting vehicle, where the focal points of ellipse E_1 are $(x, 0)$ and $z_1 = (x', 0)$, and the focal points of ellipses E_j are z_{j-1} and z_j for $j \in [2:n+1]$. Similarly, let ellipses E'_1, \dots, E'_{n+1} be the corresponding ellipses for the waiting vehicle, where the focal points of ellipse E'_1 are $(y, 0)$ and $z_1 = (x', 0)$, and the focal points of ellipse E'_j are the same as that of ellipse E_j for $j \in [2:n+1]$.

Clearly, for each point on an ellipse $E_j, j \in [2:n+1]$ the sum of the distances of the point to the focal points equals the distance between the two focal points plus w , where w is the slack of the tour. For each point on an ellipse $E'_j, j \in [2:n+1]$, this distance equals the distance between the focal points plus $w - \delta$. Because for every point on ellipse $E'_j, j \in [2:n+1]$ the sum of the distances to its foci is δ units smaller than the corresponding sum for ellipse E_j , and because E'_j and E_j have the same focal points, it follows that for every location that is covered by E'_j the distance to the nearest location that is not covered by any of the ellipses $E_h, h \in [2:n+1]$ is at least $\delta/2$ (Property 1). We need the following proposition, which is easy to show.

PROPOSITION 1. *Let E and E' be two ellipses with semi-axes that are parallel to the Cartesian axes, and let $(a, 0)$ and $(c, 0)$ be the end points of the horizontal semi-axis of E' , and $(b, 0)$, and let $(d, 0)$ be the end points of the horizontal semi-axis of E so that $a < b < c < d$. Assume that the horizontal semi-axis of E is shorter than the horizontal semi-axis of E' , i.e., $d - b < c - a$, and that the vertical semi-axis of E is properly longer than the vertical semi-axis of E' . Let C (C') be the area that is covered by E (respectively E') but is not covered by E' (respectively E), and let L be a horizontal line L that intersects E' . Then the intersection between C (C') and every horizontal line is a (possibly empty) interval, and the length of this interval is at least $d - c$ (respectively, at most $b - a$).*

Consider ellipses E_1 and E'_1 that have focal points $(x, 0), (x', 0)$ and $(y, 0), (x', 0)$, respectively. For every point on ellipse E_1 , the sum of the distances to the focal points is $w + x' - x$ where w is the slack of the tour. For E'_1 , the corresponding sum of distances is of course the same (each vehicle has the same time to reach $(x', 0)$), but the remaining slack of this tour is only $w - \delta$ due to earlier waiting. Observe that the horizontal semi-axes of both ellipses have the same length $w + x' - x$, whereas the vertical semi-axis of E_1 is longer than the corresponding semi-axis of E'_1 . Hence, ellipses E_1 and E'_1 satisfy the requirements of Proposition 1.

Let C' be the area that is covered by E'_1 but not by E_1 . Altogether, it follows that the area C'' of locations that are covered by at least one of the ellipses $E'_j, j \in [1:n+1]$ but are not covered by any of the

ellipses $E_j, j \in [1:n+1]$ is part of C' . Call C'' the critical area. Now consider a horizontal straight line L that intersects the critical area C'' . From Proposition 1, it can be derived that the intersection between L and the critical area C'' is contained in a segment S of L with length at most $\delta/2$. Note that each location of S is covered by E'_1 . In the following, we show that the intersection of L with all ellipses $E_h, h \in [1:n+1]$, which is not covered by any of the ellipses $E'_j, j \in [1:n+1]$, contains a segment of length at least $\delta/2$. There are two cases.

Case 1. All intersection points between the area that is covered by the ellipses $E'_j, j \in [1:n+1]$ and line L lie within the area that is covered by E'_1 . This does not imply that L does not intersect with the area covered by one of the ellipses E_j for $j \in [2:n+1]$. When it does, however, each such intersection point is also covered by E'_1 . Then, by Proposition 1, the intersection between line L and the area that is covered by E_1 but not covered by any of the ellipses $E'_j, j \in [1:n+1]$ contains a segment S' of length $\geq \delta/2$.

Case 2. Line L contains a point that is covered by an ellipse $E'_j, j \in [2:n+1]$ but is not covered by E'_1 . Consider such an intersection point (z_1, z_2) on L that is farthest from E'_1 , more exactly, (z_1, z_2) has the maximal distance to the nearest location (z'_1, z_2) on L that is covered by E'_1 . Without loss of generality let $z'_1 < z_1$. Then, each location (z'_1, z_2) with $z'_1 > z_1$ is not covered by any of the ellipses E'_j for $j \in [1:n+1]$. Using Property 1, we can conclude that the distance between (z_1, z_2) and the next point that is not covered by any of the ellipses E_2, \dots, E_{n+1} is at least $\delta/2$. Thus, the intersection between line L and all locations that are covered by ellipses $E_j, j \in [2:n+1]$ but are not covered by any ellipse $E'_h, h \in [2:n+1]$ contains a segment S' of length at least $\delta/2$. By our construction, we can also assume that no location of S' is covered by E'_1 .

Segments S and S' are disjoint because S is covered by E'_1 and S' contains no point that is covered by E'_1 . Because S is not longer than S' and L was an arbitrary horizontal line, it follows easily that the critical area cannot be larger than the area that is covered only by ellipses $E_j, j \in [1:n+1]$ but not by ellipses $E'_h, h \in [1:n+1]$. This is a contradiction to the assumption that waiting can be advantageous.

A similar argument holds for the case that there is a customer at location x . The main difference to the case of no customer at x is that ellipse E'_1 is replaced by two smaller ellipses that are completely encompassed by the former.

So far, it has been assumed that the time of the arrival of the new customer t_c is known. However, the optimal strategy is independent of the value of t_c , thus we can eliminate it from our consideration and conclude that even for unknown t_c , it is optimal to drive without waiting. \square

4.2. Two Vehicles

As we have shown in §4.1, for a single vehicle and a nonrestrictive service region it is optimal to drive without waiting. However, the problem becomes much more complex when several vehicles are involved, because the areas they cover are overlapping. For example, it is not optimal to have two or more vehicles return to the depot before the new customer has arrived because the areas covered by these vehicles coincide. Having at least one vehicle waiting some distance from the depot would definitely cover some additional area.

In this section, we determine the distance that two vehicles should ideally have from the depot at the time of the arrival of the new customer. Because this depends on the actual tours we study a simple scenario that provides some general insights. Therefore, we deviate from the WDP in the rest of section by allowing the vehicles to wait anywhere on their tour (not just at customer locations). Furthermore, we assume that the two vehicles approach the depot from opposite sides and that there are no intermediate customers (see Figure 2).

The following theorem is interesting because it shows that, for every amount of remaining time until the vehicles have to be back at the depot, an optimal distance from the depot can be determined. This result gives hints at how fast the vehicles should approach the depot; it is used in §5 to design a heuristic for more general situations.

THEOREM 4. *Let there be two vehicles that approach the depot from opposite sides where each vehicle has a remaining time t_p for driving or waiting, or both. Then the optimal distance from the depot is approximately $0.5059 \cdot t_p$ for each vehicle when the service region is nonrestrictive.*

PROOF. We give only a sketch of the proof by omitting some straightforward but tedious computations. Let $a = 1/2 \cdot t_p$, $a > 0$ denote the major semiaxis of the ellipses. First, we derive the size of the area covered by both vehicles depending on the remaining maximal driving distance $2a$ and the distance between the foci of both ellipses. The boundary of the upper halves of the ellipses can be described by the




Figure 2 Two Vehicles Moving from Opposite Sides Toward the Depot in the Center

following functions:

$$f_1(z, e_1) = b(e_1) * \sqrt{1 - ((z + e_1)/a)^2}$$

for $z \in [-a - e_1, a - e_1]$,

$$f_2(z, e_2) = b(e_2) * \sqrt{1 - ((z - e_2)/a)^2}$$

for $z \in [-a + e_2, a + e_2]$,

$$b(e) = \sqrt{a^2 - e^2},$$

where the depot is at $(0, 0)$, and the vehicles are at $-2e_1$ and $2e_2$, respectively. Setting $f_1(z, e_1) = f_2(z, e_2)$ leads to the intersection at $s = e_2 - e_1$. The total area $A(e_1, e_2)$ that is covered can then be calculated as

$$A(e_1, e_2) = 2 \left(\int_{z=-a-e_1}^{e_2-e_1} f_1(z, e_1) dz + \int_{z=e_2-e_1}^{a+e_2} f_2(z, e_2) dz \right)$$

$$= \frac{b(e_1)}{a} \left(e_2 \sqrt{a^2 - e_2^2} + a^2 \arcsin \frac{e_2}{a} \right)$$

$$+ a \frac{\pi}{2} (b(e_1) + b(e_2))$$

$$+ \frac{b(e_2)}{a} \left(e_1 \sqrt{a^2 - e_1^2} + a^2 \arcsin \frac{e_1}{a} \right).$$

For $0 \leq e_i \leq a_i$, $A(e_1, e_2)$ is a concave function, because the Hessian matrix is negatively semidefinite. Therefore, there is at most one local maximum, which is then also a global maximum. For $e = e_1 = e_2$ we obtain

$$A'(e) = A(e, e)$$

$$= 2 \left(\frac{b(e)}{a} \left(e \sqrt{a^2 - e^2} + a^2 \arcsin \frac{e}{a} \right) + a \frac{\pi}{2} b(e) \right).$$

Setting the first derivative to zero and solving this equation numerically yields a local maximum at $e^* \approx 1.0118 \cdot a = 0.5059 \cdot t_p$. Because $A(e_1, e_2)$ is symmetric to $e_1 = e_2$, the derivative that is perpendicular to $e_1 = e_2$ has to be zero as well, i.e., (e^*, e^*) is also the local maximum for $A(e_1, e_2)$ in general. □

The theorem states that both vehicles should seek a position such that their distance to the depot is about half the total distance they can still drive. For unknown time t_c , the ideal strategy is to move slowly toward the depot as the remaining time elapses, approximately at half of the normal speed of the vehicle. Also due to the concavity of function A that describes the size of the covered area, it follows that the optimal waiting strategy for either vehicle is to reach the changing optimal position as quickly as possible, and then to follow it toward the depot.

5. Heuristics for the General Case

Because the WDP is NP-hard, it is unlikely that polynomial time algorithms exist for generating optimal

solutions. For reasonably large problems, it is therefore common to design heuristics to find good solutions efficiently. In this paper, we use an evolutionary algorithm (EA) for finding good solutions and compare it with some simple deterministic heuristics. EAs are iterative stochastic search methods based on the principles of natural evolution. Starting with a set of candidate solutions (population), in each iteration (generation) of the algorithm new solutions (individuals) are created by the following process: Select two parents, construct a new solution based on these two old solutions (crossover), and slightly change the resulting solution (mutation). By repeatedly selecting good solutions for reproduction and creating new solutions based on the selected individuals, the solutions “evolve” and become increasingly better adapted to the problem at hand. A more detailed introduction to EAs is out of the scope of this paper; the interested reader is referred to Michalewicz (1996) and Davis (1991).

The EA used here has the following characteristics: The search space is the set of all waiting strategies. For a problem with m tours and n customers, a solution is represented by a string of $(n + 2 \cdot m)$ real values—one value for the waiting time at each customer and two additional values for each tour for the waiting times at the depot at the beginning and the end of the tour. Each value represents the fraction of the slack of the tour that is used for waiting at the particular location. Hence, the sum of the values corresponding to a tour is one. Because an application of the mutation or crossover operator can change the sum of these values, they are always normalized afterward. The real-valued representation and the normalization step allow the use of standard genetic operators. We use a linear ranking selection method, i.e., the individuals in the population are ordered and their selection probability decreases linearly from best to worst. For crossover between two parents, a random substring of the string of real values from one parent is selected and replaces the corresponding substring in the other parent to create one child (two-point crossover). An individual is mutated by adding to each value a normally distributed random value. The EA is run with a population size of 100, for 100 generations, where in each generation 99 children are produced, replacing the old population except for the best solution found so far (elitism).

The tricky part of applying an EA to the WDP is the evaluation of a solution, because for an arbitrary number of vehicles and arbitrary tours it is not clear how to calculate the probability of being able to include an additional customer into one of the tours. We chose to estimate this probability by random sampling. To evaluate a single waiting strategy during the course of the EA, 100 randomly generated customers

are checked one by one for feasibility of insertion. The percentage of customers that could be inserted then serves as an estimate of the solution quality. To ensure fair comparisons, all individuals from one generation are tested against the same set of 100 customers (see Branke 2001b). The optimization behavior of the EA was compared with the results of some simple heuristics, namely *NoWait* (which has been proven optimal for the single vehicle case) and the following five other waiting strategies:

- *Depot*. Wait at the depot as long as possible before starting, i.e., for every tour r the waiting time assigned to the depot equals the slack of the tour w_r . This is basically the opposite of the *NoWait* strategy.

- *MaxDist*. Wait at the customer location with maximal distance from the depot, i.e., for every tour r a waiting time that equals the slack of the tour w_r is assigned to the customer location with maximal distance from the depot. The idea here is keep the vehicles as far apart as possible, covering a large area on their way back to the depot.

- *Location*. Wait at each customer location for the same time, i.e., for every tour r with n_r customers each of its customers is assigned a waiting time w_r/n_r . This strategy is a kind of compromise between the first two heuristics: *NoWait* and *Depot*.

- *Distance*. Spread the whole waiting time proportionally to the distance driven, i.e., for every tour $r = (0, c_1, c_2, \dots, c_{n_r}, 0)$ customer c_i , $i \in [1:n_r]$ receives waiting time $w_r \cdot d_{c_i, c_{i-1}} / \sum_{j=1}^{n_r} d_{c_j, c_{j-1}}$. Basically, this is a variant of the previous heuristic, with distance rather than number of customers used to divide up waiting time.

- *Variable*. Drive without waiting until the time to drive the remaining distance to the depot is equal to the slack time. Then distribute the available waiting time to the remaining customers proportional to remaining driving distances. This heuristic basically corresponds to the strategy proven optimal in Theorem 4 for two vehicles with customers on a straight line.

We tested two versions of the EA: the standard EA (referred to as EA1) and a version where the simple heuristics were for seeding the initial population (EA2). In the latter case, the waiting strategies that were obtained by these six simple heuristics were put into the initial population, and the other 94 members were generated randomly in the hopes of providing the EA with some helpful information.

6. Empirical Results

In this section, we empirically compare the five simple waiting strategies and the waiting strategy found by the EA. Tests were performed on the basis of the following VRP problem instances from Beasley’s

OR library (Beasley 1990): c50, c75, c100, c100b, c120, c150, and c199. For each VRP instance, the tours of an optimal solution were used as the given tours for the WDP. The maximum allowed travel time T for an instance was defined as the time required for the longest tour.

Because we showed that the service region influences the optimal waiting strategy, we chose a service region that seems realistic for many practical problems. For each instance, the service region was defined as the rectangle that is limited by the maximum and minimum values of x and y coordinates of all given customer locations. To compare the quality of the solutions found by the different waiting strategies on a specific problem instance, 1,000 new customers (uniformly distributed within the service region) and arrival times uniformly distributed in $[0:T]$ were created. To evaluate a solution, we checked how many of those customers could have been serviced (one at a time). The results reported below are averaged for every heuristic over 20 such test runs of 1,000 customers each. For every test run, the EA was started with a different seed for the random number generator.

6.1. Ability to Insert an Additional Customer

The effect of the waiting strategy on the ability to integrate a new customer into one of the tours is shown in Table 1. The table shows the mean and the standard error for the number of customers that could *not* be successfully inserted, out of the 1,000 customers tested, and the relative performance compared to the results of *NoWait*. The average relative performance over all test problems is summarized in Table 2.

As can be seen clearly, waiting at the depot as long as possible at the beginning of a tour (*Depot*) is the worst strategy for each of the test instances. This strategy fails insertion up to 96% more frequently than the reference strategy *NoWait*. For each of the other heuristics, there are test problems where they perform quite well, and other test problems where they are less successful. Not to wait is sometimes among the best approaches (e.g., c50), but performs rather poorly on other test instances, such as on c199 or c100b. It seems that the quality of this “natural” approach depends strongly on the test problem.

Waiting at the customer farthest from the depot (*MaxDist*) on average performs slightly worse than *NoWait*, and there is no test problem where it turned out to perform best. A possible reason is that the service region does not favor this strategy because it is restricted and therefore the vehicles tend to wait near the border of the service region.

The other three simple heuristics and the EA all perform significantly better than *NoWait*. The *Variable* strategy derived from the simple two-vehicle case is

Table 1 Number of Customers Which Could *Not* Be Inserted for the Different Waiting Strategies and All Test Instances

Problem	Waiting strategy	Customers not inserted	Relative perf. (%)
c50	<i>NoWait</i>	449.2 ± 2.2	100.0
	<i>Depot</i>	882.2 ± 1.8	196.4
	<i>MaxDist</i>	514.9 ± 3.0	114.6
	<i>Location</i>	498.1 ± 3.0	110.9
	<i>Distance</i>	493.8 ± 3.0	109.9
	<i>Variable</i>	438.3 ± 2.2	97.6
	EA1	485.5 ± 3.0	108.1
	EA2	442.5 ± 2.3	98.5
c75	<i>NoWait</i>	410.6 ± 3.6	100.0
	<i>Depot</i>	594.7 ± 2.4	144.8
	<i>MaxDist</i>	392.2 ± 3.6	95.5
	<i>Location</i>	371.7 ± 3.7	90.5
	<i>Distance</i>	372.4 ± 3.5	90.7
	<i>Variable</i>	361.0 ± 3.7	87.9
	EA1	371.3 ± 3.9	90.4
	EA2	370.5 ± 4.7	90.2
c100	<i>NoWait</i>	356.7 ± 4.6	100.0
	<i>Depot</i>	592.7 ± 4.1	166.1
	<i>MaxDist</i>	415.9 ± 4.1	116.6
	<i>Location</i>	356.9 ± 4.6	100.0
	<i>Distance</i>	354.1 ± 4.1	99.3
	<i>Variable</i>	334.7 ± 4.7	93.8
	EA1	340.9 ± 5.1	95.6
	EA2	351.3 ± 4.6	98.5
c100b	<i>NoWait</i>	491.5 ± 3.2	100.0
	<i>Depot</i>	626.2 ± 4.0	127.4
	<i>MaxDist</i>	433.0 ± 3.2	88.1
	<i>Location</i>	416.1 ± 2.9	84.6
	<i>Distance</i>	432.7 ± 3.1	88.0
	<i>Variable</i>	417.7 ± 3.2	85.0
	EA1	418.8 ± 3.1	85.2
	EA2	418.8 ± 3.4	85.2
c120	<i>NoWait</i>	399.8 ± 4.6	100.0
	<i>Depot</i>	611.5 ± 4.2	152.9
	<i>MaxDist</i>	413.4 ± 4.5	103.4
	<i>Location</i>	390.6 ± 4.8	97.7
	<i>Distance</i>	397.8 ± 4.6	99.5
	<i>Variable</i>	370.6 ± 4.6	92.7
	EA1	388.0 ± 4.8	97.0
	EA2	373.3 ± 4.6	93.4
c150	<i>NoWait</i>	378.9 ± 2.2	100.0
	<i>Depot</i>	520.8 ± 3.1	137.5
	<i>MaxDist</i>	362.0 ± 3.1	95.6
	<i>Location</i>	324.5 ± 2.7	85.6
	<i>Distance</i>	330.1 ± 2.3	87.1
	<i>Variable</i>	323.0 ± 2.4	85.3
	EA1	328.6 ± 2.6	86.7
	EA2	326.8 ± 2.5	86.2
c199	<i>NoWait</i>	370.1 ± 3.7	100.0
	<i>Depot</i>	438.0 ± 3.7	118.3
	<i>MaxDist</i>	320.0 ± 3.3	86.4
	<i>Location</i>	301.8 ± 3.4	81.5
	<i>Distance</i>	312.9 ± 3.3	84.5
	<i>Variable</i>	298.2 ± 4.1	80.6
	EA1	303.0 ± 3.0	81.9
	EA2	300.6 ± 2.4	81.2

Note. For each case, the mean and the standard error, as well as the relative performance in percent with respect to *NoWait*, are reported.

Table 2 Number of Additional Customers that Could Not Be Inserted, Relative Percentage, Compared with Not Waiting At All

Waiting strategy	Relative performance (%)
<i>NoWait</i>	100.0
<i>Depot</i>	149.1
<i>MaxDist</i>	100.0
<i>Location</i>	93.0
<i>Distance</i>	94.2
<i>Variable</i>	89.0
<i>EA1</i>	92.6
<i>EA2</i>	90.0

clearly the best simple heuristic, showing that the strategy is not only optimal for two vehicles on a straight line, but also performs very well in more general cases. Distributing the waiting time equally over all customer locations (*Location*) is also a quite successful heuristic, with, on average, only 93% of the failures of the reference solution and a worst case of 110.9% for test problem c50. Waiting proportionally to the distance driven (*Distance*) performs similar and is only slightly worse, on average (94.2%).

The EA also belongs to the better heuristics: Even without seeding it is able to reduce the number of failures to, on average, 92.6% compared with the reference solution (*EA1*). The EA with seeding (*EA2*) is even better. However, surprisingly, it is not as good as the best deterministic heuristic (*Variable*), although that solution was present in the initial population. This means that the EA is not always able to recognize the value of the seed. A problem is that the quality estimate from 100 customers that was used in the EA cannot guarantee the survival of a good solution. Better sampling strategies might solve this problem as has been suggested in Branke (2001a).

Summarizing, an appropriate waiting strategy significantly increases the probability for a successful insertion of a new customer when compared with the “natural” strategy *NoWait*. The *Variable* strategy derived from the theoretical results in §4.2 performs best. Also successful are the simple heuristics to distribute the slack time of the tours equally (*Location*) or according to the length of the tour (*Distance*). The proposed EA also yielded quite successful results; however, it did not quite reach the performance of the *Variable* strategy. Nevertheless, the approach has potential for further improvements, and is probably the only heuristic that can be easily adapted to nonuniformly distributed customer locations (see §7).

6.2. Length of Detour

As we have shown in the previous section, an appropriate waiting strategy can significantly increase the probability of being able to integrate an additional customer into the given tours. However, there is an

additional advantage of waiting. Because the waiting vehicles have not yet proceeded as far on their tours at the time when the additional customer becomes known, they have more opportunities to integrate the customer at a convenient position into the sequence of customers of a tour. Thus, it can be expected that compared with *NoWait*, a smaller detour is necessary to integrate the additional customer (in fact, because some time was spent waiting, the maximal detour possible is smaller than when the vehicle was driving without waiting).

The average length of the detour to insert a customer that can be successfully integrated is given in Table 3, and the average relative performance over all test problems is displayed in Table 4. The observed savings with respect to *NoWait* are impressive. Ignoring the strategy to wait at the depot for now, all suggested waiting strategies save between 25% and 35% of the driving distance required to insert a new customer, on virtually all problem instances. The best heuristics are *Location* and the EA. The best strategy with respect to the number of inserted customers, *Variable*, only ranks fifth with respect to the required detour.

Waiting at the depot not only often fails to integrate new customers (as has been shown in §6.1), but also results in longer average detours to insert a customer. Because vehicles have different amounts of time available for waiting, some of them have used up all of their waiting time early in their tours, leaving it up to a small number of remaining vehicles to insert the new customer, which leads to longer detours.

Figure 3 compares the different heuristics with respect to the average values for both examined criteria, the number of customers not inserted, and the average length of the detour per customer inserted. The impressive improvements possible by an appropriate waiting strategy are clearly visible, with the waiting strategies found by the EA and *Variable* performing best.

7. Problem Variants

So far, we restricted the WDP to the insertion of a single new customer, and a uniform distribution of possible customer locations. In this section, we briefly discuss some aspects of the WDP when relaxing these restrictions.

7.1. Nonuniform Distribution of Customer Locations

If the probability distribution of the newly arriving customer is nonuniform and it is known that there are areas with a higher probability of customers arriving, it will, in general, make sense to wait at such areas. Obviously, the benefits of waiting depend on the particular probability density function, and instances

Table 3 Average Detour per Customer Inserted, for the Different Waiting Strategies and All Problem Instances

Problem	Waiting strategy	Average detour	Relative perf. (%)
c50	<i>NoWait</i>	5.37±0.07	100.0
	<i>Depot</i>	2.58±0.07	48.0
	<i>MaxDist</i>	4.76±0.07	88.6
	<i>Location</i>	3.59±0.04	66.9
	<i>Distance</i>	3.66±0.04	68.3
	<i>Variable</i>	5.11±0.05	95.2
	<i>EA1</i>	3.76±0.05	70.0
	<i>EA2</i>	4.94±0.08	92.1
c75	<i>NoWait</i>	9.62±0.14	100.0
	<i>Depot</i>	11.46±0.14	119.0
	<i>MaxDist</i>	7.16±0.11	74.4
	<i>Location</i>	6.25±0.10	65.0
	<i>Distance</i>	6.66±0.10	69.2
	<i>Variable</i>	6.97±0.07	72.5
	<i>EA1</i>	6.25±0.11	64.9
	<i>EA2</i>	6.27±0.12	65.2
c100	<i>NoWait</i>	9.44±0.14	100.0
	<i>Depot</i>	15.78±0.27	167.1
	<i>MaxDist</i>	7.49±0.10	79.3
	<i>Location</i>	6.93±0.11	73.3
	<i>Distance</i>	6.72±0.10	71.1
	<i>Variable</i>	7.05±0.08	74.7
	<i>EA1</i>	6.89±0.13	73.0
	<i>EA2</i>	6.88±0.11	72.8
c100b	<i>NoWait</i>	20.15±0.25	100.0
	<i>Depot</i>	14.62±0.24	72.5
	<i>MaxDist</i>	15.20±0.18	75.4
	<i>Location</i>	13.55±0.15	67.2
	<i>Distance</i>	14.35±0.15	71.2
	<i>Variable</i>	13.11±0.16	65.1
	<i>EA1</i>	13.53±0.15	67.2
	<i>EA2</i>	13.31±0.17	66.1
c120	<i>NoWait</i>	16.27±0.27	100.0
	<i>Depot</i>	21.70±0.52	133.4
	<i>MaxDist</i>	11.11±0.16	68.3
	<i>Location</i>	9.38±0.13	57.7
	<i>Distance</i>	9.49±0.15	58.3
	<i>Variable</i>	8.58±0.13	52.7
	<i>EA1</i>	9.11±0.16	56.0
	<i>EA2</i>	8.67±0.14	53.3
c150	<i>NoWait</i>	9.05±0.11	100.0
	<i>Depot</i>	11.28±0.20	124.7
	<i>MaxDist</i>	5.41±0.07	59.8
	<i>Location</i>	4.84±0.05	53.5
	<i>Distance</i>	5.03±0.05	55.6
	<i>Variable</i>	5.59±0.05	61.8
	<i>EA1</i>	4.71±0.06	52.0
	<i>EA2</i>	4.75±0.06	52.5
c199	<i>NoWait</i>	8.72±0.12	100.0
	<i>Depot</i>	12.08±0.20	138.5
	<i>MaxDist</i>	5.04±0.06	57.8
	<i>Location</i>	4.72±0.07	54.1
	<i>Distance</i>	4.91±0.07	56.4
	<i>Variable</i>	5.04±0.05	57.8
	<i>EA1</i>	4.58±0.06	52.5
	<i>EA2</i>	4.61±0.06	52.8

Note. For each case, the mean, the standard error and the relative performance in percent with respect to *NoWait* are reported.

Table 4 Relative Average Distance of Detour to Integrate an Additional Customer Compared with Not Waiting At All

Waiting strategy	Relative performance (%)
<i>NoWait</i>	100.0
<i>Depot</i>	114.7
<i>MaxDist</i>	71.9
<i>Location</i>	62.5
<i>Distance</i>	64.3
<i>Variable</i>	68.5
<i>EA1</i>	62.2
<i>EA2</i>	65.0

could be constructed where waiting is more (or less) beneficial than for the case of uniform distribution. Note that the adaptation of the EA to this problem variant is straightforward, because one can just use the appropriate (known) distribution of customer locations to generate test cases.

7.2. Two or More New Customers

Now consider the case wherein more than one new customer is arriving:

- Contrary to the case of a single new customer, it may now be beneficial to have areas covered by more than one vehicle. If several new customers fall into such an area and cannot all be serviced by a single vehicle, it may still be possible to serve them by one of the other vehicles covering that area. However, an area covered by a single vehicle is still more valuable than an area also covered by another vehicle.

- If the service region is large and the number of new customers is small, it is unlikely that more than one customer fall into an area covered by more than one vehicle. Then, our experimental results concerning the relative quality of the different waiting strategies will approximately hold.

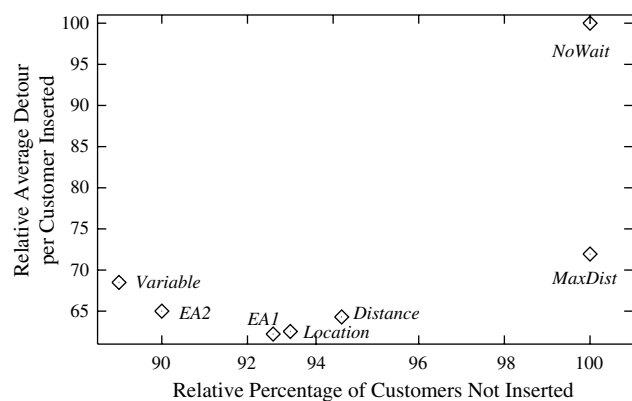


Figure 3 Comparison of the Heuristics with Respect to the Average Number of Customers Not Inserted and the Average Length of the Detour for a Successful Insertion a New Customer

Note. Heuristic *Depot* has been omitted for clarity because it performs much worse than the other strategies.

- For the case of a single vehicle it has been shown in Theorem 3 that for one new customer and a non-restrictive service region, *NoWait* is always the optimal strategy. In the following, we give an example that shows that for the same case, but with more than one new customer, waiting can be advantageous. Consider a vehicle that services its last customer at location $(-4, 0)$, which corresponds to a distance of 4 from the depot at location $(0, 0)$. The vehicle has eight time units left. It may either proceed to the depot without waiting, or it may wait for up to four time units at the last customer. Let us now assume that new customers arrive within a nonrestrictive service region after four time units. Depending on how long the vehicle waited, it will now be somewhere on its way from the last customer to the depot. To test what waiting strategy is best, we generated random test instances with one, two, and three new customers. For each of 100,000 randomly generated sets of new customers we determined the number of customers that could be inserted (assuming the best ordering of servicing the customers), depending on the waiting time. The results are depicted in Figure 4.

The results show that (in accordance with Theorem 3) for a single new customer, a vehicle that does not wait has the highest probability of successful insertion. For a larger number of two and three new customers, a smaller fraction of the new customers could be inserted (often only one out of the two or three new customers can be inserted). However, note that waiting becomes advantageous for more than one new customer. In the example, a waiting time of 0.52 is best for two new customers; for three new customers, an even larger waiting time of about 0.92 is best. These results indicate that waiting may be an even more powerful strategy when more than one (but still few) new customers have to be inserted.

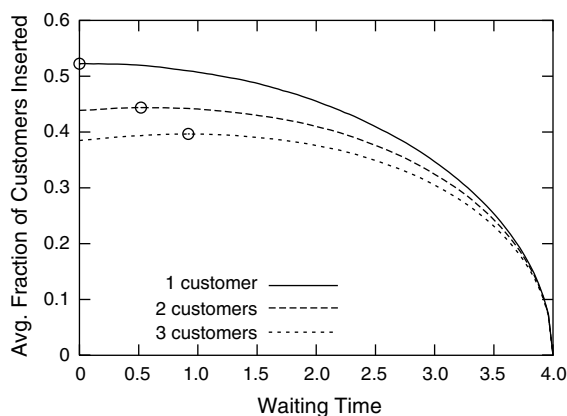


Figure 4 Average Fraction of the New Customers that Can Be Inserted into the Tour, Depending on the Time Spent for Waiting and Number of New Customers

Note. For each different number of new customers, the best waiting time is indicated by a circle.

- Each new customer that is inserted into a tour reduces the tour's slack, and thus the size of the remaining ellipses. It may then be advantageous for all other vehicles to speed up and to enlarge the ellipses closer to the depot (and probably with larger overlap to the ellipses of the vehicle that has inserted a new customer). Therefore, if several customers are arriving over time, a reactive waiting strategy may have to be considered.

- If the number of new customers becomes very large so that there will always be enough new customers for the vehicles to stay busy, waiting cannot be an advantage. With an increasing number of new customers, though, the whole concept of using a priori tours becomes questionable.

- Again, the proposed EA can easily incorporate the expected arrival of more than one customer simply by adapting the probability distribution for generating test cases.

8. Conclusion

In this paper, we have studied a dynamic VRP where a single new customer arrives at a uniformly chosen random location after the vehicles have left the depot. We examined the problem of choosing an appropriate waiting strategy that maximizes the probability for being able to insert the new customer into the otherwise fixed tours.

We have shown that the problem of finding an optimal waiting strategy is NP-complete no matter whether the time of arrival is known or unknown and even when each tour has at most a constant number of customers and tours do not cross between the first customer after the depot and the last customer before the depot. However, for the case of a single vehicle, we have shown that not to wait is the optimal strategy when the service region where the new customer may arrive is nonrestrictive (i.e., it contains all reachable locations). For simple problems with two vehicles, we showed that waiting is beneficial and derived an optimal waiting strategy. For the general problem, we empirically compared several heuristic waiting strategies as well as waiting strategies evolved by an evolutionary algorithm. Of all the heuristic waiting strategies tested, the one derived from the theoretical considerations on a simple two-vehicle case performed best. Overall, the obtained results clearly demonstrate the advantage of an appropriate waiting strategy: Compared with the reference strategy of not to wait, the best waiting strategies were able to reduce the probability that the customer cannot be serviced by approximately 10%, while the average length of the detour per successful insertion of a customer was reduced by approximately 35%. In other words, a good waiting strategy can increase the probability of being able to insert an additional customer,

at the same time reducing the average cost of inserting the customer (assuming cost depends on driving distance).

There remain several avenues for future research. First, our paper focuses on the case of inserting a single additional customer. A natural next step would be to explore further the problem variants discussed briefly in §7, namely a nonuniform distribution of new customers and the arrival of more than one new customer. Our assumption was that the order of customers on a predetermined tour is not changed and a new customer is inserted between two other customers on the tour. Allowing to change the predetermined tour to some extent and allowing to switch customers between tours are interesting aspects of the problem, as well. Finally, the presented EA may be improved, e.g., by using better sampling strategies for evaluation, as we have suggested.

Acknowledgments

The authors would like to thank the anonymous referees for their helpful comments. Also, they would like to thank the previous editor-in-chief Gilbert Laporte for pointing them to the references of Mitrovic-Minic and Laporte (2004) and Mitrovic-Minic, Krishnamurti, and Laporte (2004), which were under review at the time of first submission.

Appendix. Proof of Theorem 1

PROOF OF THEOREM 1. First consider the case that the customer arrival time is known. We reduce 3-SAT to our problem. Let $\mathcal{C} = \{C_1, C_2, \dots, C_m\}$ be a set of clauses each of size three over a set $V = \{v_1, v_2, \dots, v_n\}$ of variables. We assume that for each variable there are exactly three positive and three negated literals in the clauses (3-SAT remains NP-complete under this restriction, see Garey and Johnson 1979). We construct an instance of WDP as follows: The maximal allowed length for a tour, T , is chosen large enough so that the tours that are defined in the following are feasible. We need the following two facts (the first is easy to show and the proof is omitted). □

FACT 1. Let $S = \{S_1, S_2, \dots, S_p\}$ be a set of disjoint circles in the service region that do not contain the depot and $t, k > 0$ given integers (compare Figure 5). It is possible to construct in polynomial time a set of tours R and a convex nonrestrictive service region with the following properties: (i) a waiting strategy for R can be found in polynomial time so that the reachable area of the




Figure 5 Circles in the Set S and Depot in the Service Region




Figure 6 Construction for the Proof of Fact 2

service outside the circles is t -covered, (ii) for every circle S_i with area A_i the inner circle S'_i with area $(1 - 1/k) \cdot A_i$ of S around the center of S_i is never reachable from tours in P , and (iii) the optimal waiting strategy for all tours up to time t is to never wait.

FACT 2. Let $\mathcal{S}_1 = \{S_1, S_2, S_3\}$ and $\mathcal{S}_2 = \{S_4, S_5, S_6\}$ be two groups of three disjoint circles each with diameter $D > 0$ that lie in the service region as described in the following (see Figure 6). Let $(x_i, y_i), i \in [1:6]$ be the center of circle S_i . Then for circles in \mathcal{S}_1 , it holds that $x_2 = x_1 + ad + D/2, y_2 = y_1 + ad + D/2, x_3 = x_1 + 2ad + D, y_3 = y_1$. Analogous relations hold for circles in \mathcal{S}_2 with $a+b$ instead of a and where $x_4 > x_3 + 4(a+b)d$ and $y_4 = y_1$. Let a, b, d, t_x , and t_y be positive integers with $b > d + 2D, d > D, t_y > t_x + 12(a+b)d$. It is possible to construct in polynomial time a tour r such that (see Figure 6): (i) there exists a waiting strategy for which the circles in \mathcal{S}_1 are $(t_x + d)$ -covered, (ii) there exists a waiting strategy for which the circles in \mathcal{S}_2 are $(t_y + d)$ -covered, (iii) there does not exist a waiting strategy so that a location in a circle in \mathcal{S}_1 is $(t_x + d + D)$ -covered, and (iv) there does not exist a waiting strategy so that a location in a circle in \mathcal{S}_1 is t_x -covered and a location in a circle in \mathcal{S}_2 is t_y -covered.

PROOF OF FACT 2. Consider the tour r that is depicted in Figure 6. Tour r is constructed so that its length is $T - (1 + 2a + 2b)d - 2D$ and the distance between the depot and location x (y) on r is $t_x - 2bd$ (respectively t_y). Furthermore, there are customers at locations $x = (x_1 + ad + D/2, y_1)$ and $y = (x_4 + (a+b)d + D/2, y_4)$ and all other customers have distance more than $4(a+b)d$ from every center of a circle in $S_1 \cup S_2$. The construction implies that (i) when the vehicle waits time $2bd$ at or before x it $(t_x + d)$ -covers any location in a circle in \mathcal{S}_1 , (ii) when the vehicle does not wait before y it $(t_y + d)$ -covers every location in a circle in \mathcal{S}_2 , (iii) when the vehicle waits more than time $d + 2D$ at or before x it does not t_y -cover a location in a circle in \mathcal{S}_2 , and (iv) when the vehicle waits less than time $d + 2D < b$ at or before x it does not t_x -cover a location in a circle in \mathcal{S}_1 . Fact 2 follows easily. □

The tour that was constructed in the proof of Fact 2 is called \mathcal{S}_1 - \mathcal{S}_2 -selecting tour. Clearly, Fact 2 is also valid when the set \mathcal{S}_2 consists of only one circle S . The corresponding tour is then called a \mathcal{S}_1 - S -selecting tour.

For each variable v_i , define two groups of circles $\mathcal{S}_{i,1} = \{S_{i,1}, S_{i,2}, S_{i,3}\}$ and $\mathcal{S}_{i,2} = \{S_{i,4}, S_{i,5}, S_{i,6}\}$ where each circle has diameter D and area $A = \pi(D/2)^2$. The circles have distances from each other as shown in Figure 6 with $a = 2$ and $b = 2$. To each positive occurrence of v_i assign one circle

in $\mathcal{S}_{i,1}$, and to each negated occurrence assign one circle in $\mathcal{S}_{i,2}$. For each clause C_j let S_j be a circle with diameter D that has a minimal distance of $13d$ from all other circles. Set $\mathcal{S} = \bigcup_{i \in [1:n], k \in [1,2]} \mathcal{S}_{i,k} \cup \{S_j \mid j \in [1:m]\}$. Let $t_x, t_y, t_z > 0$ with $t_y > t_x + 12(a+b)d$ and $t_z > t_y + 12(a+b)d$. For each circle S in $\mathcal{S}_{i,1}$ ($\mathcal{S}_{i,2}$) define $t_S = t_x$ (respectively $t_S = t_y$), $i \in [1:n]$. For each circle $S = S_j$, $j \in [1:m]$ define $t_S = t_z$.

Before we define the details we describe the general idea of the construction. There are two sets of tours R_1 and R_2 . Set R_1 reflects the occurrence of variables in the clauses and allows to model a truth assignment for the variables (by defining a suitable waiting strategy). R_1 consists of $\mathcal{S}_{i,1}$ - $\mathcal{S}_{i,2}$ -selecting ($\mathcal{S}_{i,k}$ - \mathcal{S}_j -selecting) tours that can cover only circles (which one depends on the waiting strategy) in one of the sets $\mathcal{S}_{i,1}$ or $\mathcal{S}_{i,2}$ (respectively $\mathcal{S}_{i,k}$ or \mathcal{S}_j). The set R_2 consists of tours that cover the whole service region for most of the time with the exception of the circles in \mathcal{S} that are only covered as long as they cannot be covered by a tour in R_1 . The question is then whether there exists a waiting strategy so that the tours in R_1 cover all circles in \mathcal{S} during a certain time interval (when they are not covered by tours in R_2).

Set R_1 consists of two types of tours. For each variable v_i let r_i be an $\mathcal{S}_{i,1}$ - $\mathcal{S}_{i,2}$ -selecting tour as in Figure 6, i.e., with $a = 2$ and $b = 2$. Furthermore, these tours are defined so that the length of the tour from the depot until x (see Figure 7) is $t_x - 4d$, from the depot until y it is t_y , and the total length is $T - 9d - 2D$. For each positive (negative) occurrence of a variable v_i in a clause C_j define the $\mathcal{S}_{i,k}$ - \mathcal{S}_j -selecting tour where $S_{i,k}$ is the circle that is assigned to the occurrence of v_i in C_j and so that $a = 2$ and $b = 4$ (respectively, $a = 4$ and $b = 2$). Furthermore, these tours are defined so that the length of the tour from the depot until x (see Figure 7) is $t_x - 8d$ (respectively, $t_y - 4d$), from the depot until y it

is t_z , and the total length is $T - 13d - 2D$. The $\mathcal{S}_{i,1}$ - $\mathcal{S}_{i,2}$ -selecting and $\mathcal{S}_{i,k}$ - \mathcal{S}_j -selecting tours are constructed so that no location of any cycle not in $\mathcal{S}_{i,1}$ (respectively $\mathcal{S}_{i,2}$) can be reached. Let R_1 be the set of the selecting tours.

The tours in set R_2 are constructed according to Fact 1 with set of circles \mathcal{S} and time parameter $t = t_z + 13d$, $k = 2(6n + m)$. Hence, the tours in R_2 ($t_z + 13d$)-cover the whole service region besides the inner circles of the circles in \mathcal{S} . Moreover, R_2 is defined so that after time $t_z + 13d$ none of the tours in P_1 cover a location that is not covered by at least one of the tours in R_2 . For each circle $S \in \mathcal{S}$ we add an additional tour r_S to set R_2 that covers the circle S until time t_S and then the vehicle has to drive away from the circle and back to the depot. Hence it covers no location in S after time $t_S + D$. Moreover, the tour r_S never covers a location in another circle $S' \in \mathcal{S}$, $S' \neq S$. The tours in $R = R_1 \cup R_2$ form the instance of WDP.

Our construction implies that every $\mathcal{S}_{i,1}$ - $\mathcal{S}_{i,2}$ -selecting ($\mathcal{S}_{i,k}$ - \mathcal{S}_j -selecting) tour can cover only circles (which one depends on the waiting strategy) in one of the sets $\mathcal{S}_{i,1}$ or $\mathcal{S}_{i,2}$ (respectively, $\mathcal{S}_{i,k}$ or \mathcal{S}_j). If a tour covers such a circle $S \in \mathcal{S}$ it can cover it (completely) in the time interval $[t_S : t_S + d]$ when it is not covered by its tour r_S . After time $[t_S : t_S + d + D]$ no location in S is covered. To maximize the probability to serve the new customer it remains to maximize the integral of the area of circles in S that is only covered by tours in R_1 over times after the circles where covered by tours in R_2 .

On the one hand, a single tour can cover (assuming the right waiting strategy) a circle S so that the integral over time and additional (with respect to tours in R_2) covered area is $\geq (d - D)(1 - 1/(2 \cdot (6n + m))) \cdot A$ because at least the inner circle is covered additionally during $[t_S + D, t_S + d]$. On the other hand, the integral over time and additional covered area of a circle in \mathcal{S} is at most $(d + D) \cdot A$ because locations on the circle can be covered additionally only during $[t_S : t_S + d + D]$. So when all $6n + m$ circles are covered, the integral over time and the additional covered area is at least $(6n + m)(d - D)(1 - 1/(2 \cdot (6n + m))) \cdot A$. Assume d was chosen such that $d > 4(6n + m)D$. When at least one circle is not covered additionally by a tour in R_1 , the integral over time and additional covered area is at most

$$\begin{aligned} & (6n + m - 1)(d + D) \cdot A \\ & < (6n + m) \cdot dA + (6n + m) \cdot DA - dA \\ & < (6n + m)dA - (6n + m) \cdot DA - \frac{1}{2} \cdot dA + \frac{1}{2}DA \\ & = (6n + m)(d - D) \left(1 - \frac{1}{2 \cdot (6n + m)} \right) \cdot A. \end{aligned}$$

Now it is sufficient to show that all circles in S can be covered additionally by tours in R_1 if and only if there exists a \mathcal{E} -satisfying truth assignment for V . Let a \mathcal{E} -satisfying truth assignment for V be given. If variable v_i is true, choose a waiting strategy for the $\mathcal{S}_{i,1}$ - $\mathcal{S}_{i,2}$ -selecting tour r_i so that $\mathcal{S}_{i,1}$ is covered in time interval $[t_{S_{i,1}} : t_{S_{i,1}} + d]$. Otherwise, assign a waiting strategy so that $\mathcal{S}_{i,2}$ is covered during $[t_{S_{i,2}} : t_{S_{i,2}} + d]$. For each $\mathcal{S}_{i,k}$ - \mathcal{S}_j -selecting tour $r_{i,k}$, assign a waiting strategy so that S_j is covered during $[t_{S_j} : t_{S_j} + d]$ when the corresponding literal is true. Otherwise, a waiting strategy is chosen so that $S_{i,k}$ is covered during $[t_{S_{i,k}} : t_{S_{i,k}} + d]$. Clearly




Figure 7 The Three Types of Tours in the WDP Instance: $i \in [1:n]$, $j \in [1:m]$

every circle of the form $S_{i,k}$ is covered in $[t_{S_{i,k}} : t_{S_{i,k}} + d]$ by some tour. Because for each clause C_j there exists a true literal, there is a path with a waiting strategy so that the circle S_j is covered during $[t_{S_j} : t_{S_j} + d]$ by some $S_{i,k}$ - S_j -selecting tour. In the other direction, there exists a \mathcal{C} -satisfying truth assignment when there exists a waiting strategy for R such that all circles in S can be covered additionally; this is easy to show.

Now consider the case that the customer arrival time is unknown. The proof is similar to that of a known arrival time. The main changes necessary for the construction are to remove for each circle in $S \in \mathcal{S}$ the tour r_S from the R_2 and to set $t_c = t_x$. \square

References

- Beasley, J. E. 1990. Or-library: Distributing test problems by electronic mail. *J. Oper. Res. Soc.* **41**(11), <http://www.ms.ic.ac.uk/info.html>.
- Bertsimas, D. J., D. Simchi-Levi. 1996. A new generation of vehicle routing research: Robust algorithms, addressing uncertainty. *Oper. Res.* **44**(2) 286–304.
- Bertsimas, D. J., G. van Ryzin. 1991. A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Oper. Res.* **39**(4) 601–615.
- Bertsimas, D. J., G. van Ryzin. 1993. Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles. *Oper. Res.* **41**(1) 60–76.
- Bianchi, L. 2000. Notes on dynamic vehicle routing—The state of the art. Technical Report IDSIA-05-01, IDSIA, Manno-Lugano, Switzerland, <http://www.idsia.ch/~leo/papers/IDSIA-05-01.pdf>.
- Branke, J. 2001a. *Evolutionary Optimization in Dynamic Environments*. Kluwer Academic, Boston, MA.
- Branke, J. 2001b. Reducing the sampling variance when searching for robust solutions. L. Spector, W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela, R. E. Smith, eds. *Genetic and Evolutionary Computation Conference (GECCO'01)*, Morgan Kaufmann, San Francisco, CA, 235–242.
- Branke, J., D. Mattfeld. 2000. Anticipation in dynamic optimization: The scheduling case. M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J. J. Merelo, H.-P. Schwefel, eds. *Parallel Problem Solving from Nature (PPSN VI)*, Vol. 1917. LNCS, Springer, Berlin/Heidelberg, Germany, 253–262.
- Davis, L. 1991. *Handbook of Genetic Algorithms*. Van Nostrand Reinhold, New York.
- Garey, M. R., D. S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York.
- Gendreau, M., J.-Y. Potvin. 1998. Dynamic vehicle routing and dispatching. T. Crainic, G. Laporte, eds. *Fleet Management and Logistics*. Kluwer, London, UK, 115–126.
- Gendreau, M., F. Guertin, J.-Y. Potvin, E. Taillard. 1999. Parallel tabu search for real-time vehicle routing and dispatching. *Transportation Sci.* **33**(4) 381–390.
- Gendreau, M., P. Badeau, F. Guertin, J.-Y. Potvin, E. D. Taillard. 1996. A solution procedure for real-time routing and dispatching of commercial vehicles. Technical Report CRT-96-24, Centre de recherche sur les transports, Université de Montréal, Montréal, Québec.
- Ichoua, S., M. Gendreau, J.-Y. Potvin. 2000. Diversion issues in real-time vehicle dispatching. *Transportation Sci.* **34**(4) 426–438.
- Kilby, P., P. Prosser, P. Shaw. 1998. Dynamic VRPs: A study of scenarios. Technical Report APES-06-1998, School of Computer Science, University of St. Andrews, St. Andrews, Scotland.
- Lichtenstein, D. 1982. Planar formulae and their uses. *SIAM J. Comput.* **11** 329–342.
- Michalewicz, Z. 1996. *Genetic Algorithms + Data Structures = Evolution Programs*, 3rd ed. Springer Verlag, Berlin/Heidelberg, Germany.
- Mitrovic-Minic, S., G. Laporte. 2004. Waiting strategies for the dynamic pickup and delivery problem with time windows. *Transportation Res. B* **38** 635–655.
- Mitrovic-Minic, S., R. Krishnamurti, G. Laporte. 2004. Double-horizon based heuristics for the dynamic pickup and delivery problem with time windows. *Transportation Res. B* **38** 669–685.
- Papastavrou, J. D. 1996. A stochastic and dynamic routing policy using branching processes with state dependent immigration. *Eur. J. Oper. Res.* **95** 167–177.
- Powell, W. B. 1986. A stochastic model of the dynamic vehicle allocation problem. *Transportation Sci.* **20**(2) 117–129.
- Powell, W. B. 1996. A stochastic formulation of the dynamic assignment problem, with an application to truckload motor carriers. *Transportation Sci.* **30**(3) 195–219.
- Powell, W. B., M. T. Towns, A. Marar. 1998. On the value of optimal myopic solutions for dynamic routing and scheduling problems in the presence of user noncompliance. Technical Report SOR-97-15, Department of Civil Engineering and Operations Research, Princeton University, Princeton, NJ.
- Psaraftis, H. N. 1988. Dynamic vehicle routing problems. B. L. Golden, A. A. Assad, eds. *Vehicle Routing: Methods and Studies*, Vol. 16. *Studies in Management Science and Systems*. North-Holland, Amsterdam, The Netherlands, 223–248.
- Psaraftis, H. N. 1995. Dynamic vehicle routing: Status and prospects. *Ann. Oper. Res.* **61** 143–164.
- Regan, A. C., H. S. Mahmassani, P. Jaillet. 1995a. Dynamic decision making for commercial vehicle operations using real-time information. *Transportation Res. Record* **1537** 91–97.
- Regan, A. C., H. S. Mahmassani, P. Jaillet. 1995b. Evaluation of dynamic fleet management systems: A simulation framework. *Transportation Res. Record* **1645** 176–184.
- Regan, A. C., H. S. Mahmassani, P. Jaillet. 1995c. Improving efficiency of commercial vehicle operations using real-time information: Potential uses and assignment strategies. *Transportation Res. Record* **1493** 188–198.
- Swihart, M. R., J. D. Papastavrou. 1999. A stochastic and dynamic model for the single-vehicle pick-up and delivery problem. *Eur. J. Oper. Res.* **114** 447–464.
- Thomas, B. W., C. C. White. 2004. Anticipatory route selection. *Transportation Sci.* **38**(4) 473–487.
- Yang, J., P. Jaillet, H. S. Mahmassani. 1999. On-line algorithms for truck fleet assignment and scheduling under real-time information. *Transportation Res. Record* **1667** 107–113.