

EVALUATING IDLE TIME POLICIES FOR REAL-TIME ROUTING OF A SERVICE VEHICLE

Complete Research

Meisel, Stephan, University of Muenster, Muenster, Germany, stephan.meisel@uni-muenster.de

Wölck, Martin, University of Muenster, Muenster, Germany, martin.woelck@uni-muenster.de

Abstract

We empirically evaluate a number of alternative decision rules for real-time routing of a service vehicle. The particular feature of the decision rules is that they aim at more efficient routing of the vehicle by deliberately introducing vehicle idle times. Our goal is to provide both routing information system providers and supply chain operators with criteria for application of these policies in practice.

The main characteristic of the class of routing problems considered is that part of the customer requests are uncertain, i.e., that part of the requests are not known in advance, but occur only while the service vehicle already is en route.

In order to derive criteria for successful application of the decision rules in practice, we compare the rules by means of computational experiments with respect to a broad variety of instances of the considered class of routing problems. In particular we derive guidelines by varying service region, time horizon and customer request behavior. Eventually, we assess the rules' performance by comparing them with ideal a posteriori solutions.

Keywords: dynamic vehicle routing, uncertain service requests, decision rules, mathematical optimization.

1 Introduction

Service vehicles play an important role in many supply chains. Last mile deliveries, recycling services and less-than-truckload trucking are only a few examples of supply chain operations that cannot be realized without deliberate use of service vehicles. However, usually careful planning is required in order to ensure that service vehicles are used efficiently. In fact, efficient vehicle routing turns out to be one of the keys to high supply chain performance.

As a consequence of both the need for efficiency and the complexity of vehicle routing problems, companies rely on information systems that process data on vehicles, on infrastructure and on service tasks for providing the decision maker with efficient routing plans.

Although information system support has been crucial already since many decades ago, its significance has increased tremendously over the past few years. Enabling technologies such as mobile communication, the Global Positioning System and big data now allow for the design and use of information systems that support real-time routing of service vehicles. In contrast to more traditional settings where routing plans are made once before a service vehicle leaves its depot, modern supply chains rely on systems that provide real-time information allowing for plan revisions while the vehicle is en route.

The ability to revise routing plans in real-time is of particular importance in the presence of uncertainties. If the information needed for planning the route of a service vehicle is revealed only gradually over time, decision makers need the opportunity to adapt routing plans to the most recent information state. New customer requests, changing traffic conditions and vehicle break-downs are only a few of many examples of why the information state might be changing over time.

In order to enable supply chain managers to make good routing decisions in uncertain environments they must have access to information systems that not only gather information about customers, traffic and vehicles, but that additionally incorporate high-quality decision rules for informing the decision maker about how routing plans should be revised.

In the literature, e.g., in Meisel (2011), such decision rules are referred to as "policies" for solving dynamic vehicle routing problems. As with all dynamic planning problems, many different classes of policies can be considered (Powell, 2011, ch. 6), and the policy that a decision maker should rely on generally depends on the type of routing problem at hand.

Recently, a number of authors (e.g., Branke et al., 2005; Meisel, 2011; Thomas, 2007) proposed idle time policies for dynamic vehicle routing. The key insight at the bottom of this type of policy is that in many real-time vehicle routing problems it is more efficient to let vehicles wait for a certain amount of time at certain locations instead of letting them move on to the next service request as quickly as possible. An idle time policy informs the decision maker based on the current state of information about when and where idle time should be allocated.

Although it has been shown that idle time policies may increase efficiency, and although these policies are usually easily applicable, they have not yet become a standard tool in supply chain information systems. In particular the fact that allocation of idle time may at first glance seem to be counterintuitive, raises the need for reliable criteria for the use of idle time policies.

In this paper we identify such criteria for a particular class of routing problems. We consider a number of different idle time policies and evaluate them empirically for the case of routing a service vehicle with stochastic customer requests. Our analysis includes both relative comparison of the individual policies and comparison with posteriori optimal policies. The model we use for determining posteriori optimal policies is new. In contrast to existing works about idle time policies, we do not aim at showing that idle time policies are beneficial for specific problem instances with carefully selected attribute values. Instead our empirical evaluation identifies attributes of problem instances that serve as criteria for making the decision about which policy to apply. We believe that identification of such criteria is an important contribution to establishing idle time policies in supply chain information systems.

The remainder of the paper is organized as follows. In Section 2 we present an overview of related work on idle time policies for dynamic vehicle routing. A detailed description of the class of dynamic vehicle

routing problems we consider in this work is given in Section 3. Section 4 contains the idle time policies that we include in our analysis and in Section 5 we propose the model used for calculating posteriori optimal solutions. Our empirical evaluation of idle time policies is presented in Section 6. Section 7 concludes the paper.

2 Related Work

The field of dynamic vehicle routing comprises all vehicle routing problems under uncertainty where the routing plans of vehicles are revised at least once while the vehicles already are en route. The research on dynamic vehicle routing is manifold. In the following paragraphs we summarize approaches that take into account idle time policies in the presence of stochastic customer requests. We refer to Pillac et al. (2013) for a general overview of approaches to dynamic vehicle routing.

An early approach featuring the allocation of idle time is due to Gendreau et al. (1999). They consider the problem of routing a fleet of vehicles with stochastic requests and soft time windows. If a vehicle's current route leads to an early arrival at the next customer location to be visited, the vehicle must wait instead of moving on immediately. Thus the vehicle movement is performed at the latest possible time in order to allow for last minute changes of the planned route. Ichoua et al. (2000) as well as Chen and Xu (2006) extend the approach of Gendreau et al. (1999) but do not propose a different way of allocating idle time. Ichoua et al. (2006) propose an extension including a threshold-based idle time policy. A vehicle that could wait at its current location is allowed to do so if the probability of a new request in the vehicle neighborhood within a certain period of time is greater or equal to a predefined threshold value. In this case idle time is allocated subject to a predefined upper bound and subject to a parameter representing some tolerance to additional lateness.

Larsen et al. (2004) consider dynamic routing of a single vehicle with late requests and soft time windows. The vehicle waits if its route will result in an early arrival at the next customer location. The available waiting time may be spent either at the current location or at one of a number of predefined idle points. First, an idle point is selected according to one of three proposed rules. Then, the decision on whether or not to move to the idle point is based on the probability of receiving at least one new request close to the idle point within the available waiting time. The vehicle moves to the idle point in case this probability exceeds a predefined threshold value. Bent and van Hentenryck (2007) apply idle time policies to a routing problem that is similar to the one considered by Larsen et al. (2004).

Branke et al. (2005) aim at maximizing the probability of being able to serve a single late request that appears at a uniformly distributed random location within a convex service region. The authors prove the problem of finding an optimal idle time policy to be NP-complete, yet derive an optimal policy for a simplified 2-vehicle problem. Moreover, they compare seven different idle time policies empirically.

Mitrovic-Minic and Laporte (2004) consider idle time policies for a dynamic pickup and delivery problem with time windows and stochastic requests. In a follow-up work on the same problem, Mitrovic-Minic et al. (2004) improve their approach by taking into account long term effects of routing decisions in the objective function.

Idle time allocation has also been considered as an option in the early works on the dynamic traveling repairman problem. The dispatching rules proposed by Bertsimas and van Ryzin (1991) as well as by Papastavrou (1996) frequently send the vehicle to an idle location in the center of the service region. As a consequence, the expected travel time for serving the next request is minimized if customer locations are uniformly distributed in the service region. Similar approaches to extended problem formulations are provided by Bertsimas and van Ryzin (1993) and by Swihart and Papastavrou (1999).

Both Thomas (2007) and Meisel (2011) present empirical comparisons of a variety of idle time policies for routing a service vehicle. Among all works discussed in this section, these two consider routing problems that are most similar to the problem class proposed in Section 3. A single service vehicle is routed in the course of a specific period of time within which new customer requests may occur at any point in time. The decision maker has to decide on which request to accept and which to reject such that the number

of customers served over the entire time horizon is maximized. Both Thomas (2007) and Meisel (2011) show that application of idle time policies can be beneficial.

However, both authors (as all authors mentioned above) experiment with selected problem instances featuring certain properties such as specific time horizon lengths. Their empirical results show which of the specific instances idle time policies should be considered for, but they do not identify more general criteria for the application of idle time policies.

3 Problem Formulation

We consider a class of routing problems that have applications in many real-world supply chains. The problems comprise a service vehicle and a set $\mathcal{S} = \{1, 2, \dots, N\}$ of geographic locations i with $i = 1$ representing the start depot, $i = N$ representing the final destination of the vehicle and $\mathcal{S} \setminus \{1, N\}$ representing customer locations. Travel distances d_{ij} between any pair (i, j) of locations are assumed to be both known and deterministic. For the sake of simplicity one unit of travel distance is assumed to correspond to one unit of time.

The vehicle is routed over the course of T time units, starting at location 1 at $t = 0$ and having to be at location N at time $t \leq T$. The set of customers is divided into two complementary parts \mathcal{S}^a and \mathcal{S}^d . \mathcal{S}^a represents customers that have requested for service in advance, i.e., before $t = 0$, and that must be visited once before time $t = T$. In contrast \mathcal{S}^d represents dynamic customers, each of which may issue at most one service request within the given time horizon. Due to the limited amount of time available, it may happen that a service request of a dynamic customer has to be rejected. However, once a dynamic service request is not rejected but confirmed, the issuing customer must be visited before $t = T$. The overall goal of the decision maker is to serve as many service requests as possible within the given time horizon.

With these definitions, the routing problem may be specified as a Markovian dynamic decision process.

3.1 Information State

At each point in time t the information to be taken into account for making plan revisions comprises both, the current state of the vehicle and the current request state of each customer. We define the request state r_{ti} of customer i at time t as

- $r_{ti} = 0$ if i has not requested for service yet
- $r_{ti} = 1$ if i has issued a request but has not been confirmed or rejected yet
- $r_{ti} = 2$ if the request of i has been confirmed
- $r_{ti} = 3$ if the request of i has been either rejected or served.

At time $t = 0$, $\forall i \in \mathcal{S}^a : r_{ti} = 2$ and $\forall i \in \mathcal{S}^d : r_{ti} = 0$. Moreover, the vehicle state at t is characterized by the vehicle's current destination $m_t \in \{1, 2, \dots, N\}$ and the remaining time δ_t that the vehicle still needs for arriving at m_t , if the vehicle visits customer i at time t , we let $m_t = i$ and $\delta_t = 0$.

Provided that we always process confirmed service requests in the order given by the shortest route comprising all customers with $r_{ti} = 2$ as well as m_t and N , the overall state S_t of the system at time t is defined as $S_t = (r_{t2}, r_{t3}, \dots, r_{t,N-1}, m_t, \delta_t)$.

3.2 Plan Revisions

Making a plan revision at a point in time t may imply resetting the vehicle destination as well as confirming and rejecting new customer requests. Note that in practice rejection of a request typically implies postponement of that request to the next day or to the next work shift. In case of a service provider operating a fleet of vehicles a rejected request may alternatively be shifted to another vehicle. For the driver's sake revisions of the vehicle's current destination only take place at times when the vehicle has reached a customer location.

Let $\mathcal{R}_t = \{i \in \mathcal{I} \mid r_{ti} = 1\}$ and $\mathcal{C}_t = \{i \in \mathcal{I} \mid r_{ti} = 2\}$. Then we refer to $\mathcal{A}_t = \mathcal{R}_t \cup \mathcal{C}_t$ as the set of active customers. A plan revision x_t at time t may consist of both confirmation (respectively rejection) decisions $x_t^c = (x_{ti}^c)_{i \in \mathcal{A}_t}$ and of a vehicle movement decision x_t^m , i.e., $x_t = (x_t^c, x_t^m)$. Such a revision is made whenever the vehicle arrives at its current destination, i.e., whenever $\delta_t = 0$, as well as whenever one or more new customer requests occur.

Let $x_{tij}^p \in \{0, 1\}$ represent the binary decision on whether or not the current routing plan includes travelling from location i to location j . Then the set of feasible plan revisions at time t is defined by the following equations:

$$\begin{aligned} x_{ti}^c &= 1 & \forall i \in \mathcal{C}_t \cup \{m_t, N\}, & \quad (1a) \\ \sum_{(i,j) \in \phi^+(i)} x_{tij}^p &= x_{ti}^c & \forall i \in \mathcal{A}_t \cup \{m_t\}, & \quad (1b) \\ \sum_{(k,i) \in \phi^-(i)} x_{tki}^p &= x_{ti}^c & \forall i \in \mathcal{A}_t \cup \{N\}, & \quad (1c) \\ \sum_{i \in \mathcal{A}_t \cup \{m_t\}} \sum_{j \in \mathcal{I}} d_{ij} x_{tij}^p &\leq T - t - \delta_t, & & \quad (1d) \\ \sum_{(i,j) \in \mathcal{E}(\mathcal{S})} x_{tij}^p &\leq |\mathcal{S}| - 1 & \forall \mathcal{S} \subseteq \mathcal{A}_t \cup \{m_t, N\}, & \quad (1e) \\ x_{tij}^p, y_i &\in \{0, 1\} & \forall i, j \in \mathcal{A}_t \cup \{m_t, N\}. & \quad (1f) \end{aligned}$$

Equation 1a ensures that each location that has previously been part of the routing plan and that has not been visited yet is included in the routing plan at time t . Equation 1b ensures that the vehicle leaves each location i (except the final destination N) that is part of the routing plan once via an edge that is part of the set $\phi^+(i)$ of all edges with origin i . Respectively, Equation 1c ensures that the vehicle arrives at each location i (except the current destination m_t) that is part of the routing plan once via an edge that is part of the set $\phi^-(i)$ of all edges with destination i . Equation 1d makes sure that the routing plan is feasible with respect to the time remaining, and Equation 1e represents the well know subtour elimination constraints with $\mathcal{E}(\mathcal{S})$ denoting the set of edges connecting any pair of locations that is part of \mathcal{S} .

3.3 State Transition and Objective

A transition from an arbitrary state S_t to a successor state S_{t+1} is fully determined by both decision x_t and the customer behavior in terms of newly occurring requests. As any request is perceived as a random influence occurring independently of x_t , each state transition can be split into two parts. The impact of decision x_t is known instantaneously at t , whereas the full impact of the random influence can only be observed at the next decision time $t + 1$.

Obviously decision x_t has no influence on δ_t , but it effects a transition from state S_t to post-decision state $S_t^x = (r_{t2}^x, r_{t3}^x, \dots, r_{t,N-1}^x, m_t^x, \delta_t^x)$, where $\delta_t^x = \delta_t$, as follows:

$$r_{ti}^x = \begin{cases} 2 & \text{if } r_{ti} = 1 \wedge x_{ti}^c = 1 \\ 3 & \text{if } (r_{ti} = 1 \wedge x_{ti}^c = 0) \vee x_{ti}^m = i \\ r_{ti} & \text{otherwise.} \end{cases} \quad (2)$$

$$m_t^x = \begin{cases} m_t & \text{if } \delta_t > 0 \\ i & \text{if } \delta_t = 0 \wedge x_{m_t i}^p = 1. \end{cases} \quad (3)$$

Moreover, representing a request of customer i in period t as a random variable $W_{ti} \in \{0, 1\}$ allows for formulation of a transition from S_t^x to S_{t+1} as

$$r_{t+1,i} = \begin{cases} 1 & \text{if } r_{ti}^x = 0 \wedge W_{t+1i} = 1 \\ r_{ti}^x & \text{otherwise,} \end{cases} \quad (4)$$

and $m_{t+1} = m_t^x$ as well as $\delta_{t+1} = \delta_t^x - 1$.

Given a particular policy $X^\pi(S_t) = (x_t^c, x_t^m)$ the number of confirmed requests at time t is denoted as

$$C_t(X^\pi(S_t)) = \sum_{i \in \mathcal{R}_t} x_{ti}^c. \quad (5)$$

Consequently the expected number of confirmations resulting from application of policy $X^\pi(S_t)$ may be defined as

$$C^\pi = \mathbb{E} \sum_{t \in T} C_t(X^\pi(S_t)). \quad (6)$$

Most supply chain managers would of course prefer a policy $X^\pi(S_t)$ that maximizes the expected number of confirmations. However, finding a good, not even optimal, policy usually is extremely hard. (For a thorough discussion of the computational challenges involved in finding a good policy, we refer to Meisel (2011)). As a consequence, most information systems supporting vehicle routing only include basic policies. Idle time policies are a straightforward, yet powerful extension of basic policies. The following Section 4 derives a number of idle time policies from a given basic policy.

4 Idle Time Policies

In Section 4.1 we specify a basic policy that serves as a starting point for deriving a number of idle time policies in Section 4.2.

4.1 Basic Policy

The basic policy pursues the principle of confirming as many new customer requests as possible at each single point in time. Additionally the policy aims at moving on to the next customer location as quickly as possible. In case of no more customer requests to be served, the vehicle moves on to the final destination, where it arrives at a point in time $t \leq T$. At a given state S_t decision $X^\pi(S_t) = x_t = (x_t^c, x_t^m)$ is derived according to the following two steps:

1. **Confirmation decisions:** In order to maximize the number of confirmations at time t , we solve an optimization problem with objective $\max C_t(X^\pi(S_t))$ subject to constraints 1a - 1f.
2. **Movement decision:** If the vehicle currently is on its way to the next location m_t , i.e., if $\delta_t > 0$, we make movement decision $x_t^m = m_t$. If the vehicle is located at a customer at time t and if there is at least one customer i with $r_{ti}^x = 2$, we solve an open Traveling Salesman Problem (TSP) with start location m_t , end location N and all customers with $r_{ti}^x = 2$. The vehicle then moves on to the next customer location according to the resulting routing plan. Note that the solution of the Traveling Salesman Problem now defines decision $x_{m_t i}^p$ in Equation 3. If no more confirmed customers are present at time t , the vehicle moves on to the final destination, i.e., $x_t^m = N$

The basic policy implies that the service vehicle may return to the final destination before the end of the driver's working shift. Considering extreme cases helps unveiling the fact that the basic policy will therefore not be considered in practice. For example, if not a single customer has requested for service at $t = 0$, following the basic policy means moving from start location to final destination immediately. This example shows that allocating idle time in order to make vehicles wait for additional request should

generally be an option. The following Section 4.2 derives a number of policies that take this option into account.

4.2 Advanced Policies

We consider three alternative idle time policies that have been proposed in the literature (see e.g., Meisel, 2011). Each of the policies makes confirmation decisions in the same way as described with the basic policy in Section 4.1. The policies differ from each other merely with respect to how decisions about movement of the service vehicle are made. At a state S_t each of the policies explicitly takes into account the available idle time τ_t resulting from the length of the currently planned route and from the time still needed for arriving at m_t . Let x_t^p represent the currently planned route at time t and let $l(x_t^p)$ denote the length of x_t^p . Then the available idle time at t results as $\tau_t = T - t - l(x_t^p) - \delta_t$. The three policies "Wait First", "Drive First" and "Distributed Wait" actively take into account τ_t for moving the vehicle as follows:

- **Wait First (WF):** The idea behind WF is to allocate as much idle time as possible at the vehicle's current location. Note that the routing behavior resulting from this approach strongly depends on whether or not an open TSP is solved in order to determine x_t^p . Using TSP solutions as routing plans implies that the vehicle is going to wait for as long as possible at the start location $i = 1$ while confirming newly occurring customer requests. As more requests are being confirmed and as time passes by the available idle time gets smaller. Only as soon as $\tau_t = 0$ the vehicle leaves the start location and visits one customer location after the other without more idle time being allocated and without new requests being accepted. In contrast to the basic policy, WF has the option of making movement decision $x_t^m = m_t$ even in case of $\delta_t = 0$. As decisions are made at each point in time, we assume that one decision may allocate only one full time period as idle time.
- **Drive First (DF):** The idea behind DF is exactly opposite to the idea behind WF. DF may only allocate idle time if the vehicle currently is located at the last customer location in x_t^p , i.e., it allocates as much idle time as possible to the customer location that is the one to be visited right before moving on to the final destination $i = N$. As a consequence the vehicle only waits if it has reached this location and if $\tau_t > 0$. The waiting process stops as soon as a new customer request occurs and gets confirmed or as soon as $\tau_t = 0$. Note that the last customer location to be visited right before moving on to the final destination may change as new customer requests occur in the course of time.
- **Distributed Wait (DW):** The DW policy aims at distributing idle time as equally as possible among the current location of the vehicle and the customer locations that still need to be visited. As with WF and DF no idle time is allocated only if $\tau_t = 0$. In case of the vehicle currently being located at a customer location or at the start location, one unit of idle time is allocated at this current location if both idle time is available, i.e., $\tau_t > 0$ and the vehicle has so far waited less than $\lceil \frac{\tau_t}{|\mathcal{C}_t|+1} \rceil$ units of time at the current location, where $|\mathcal{C}_t| + 1$ is the number of locations that the current route x_t^p contains if the final destination is ignored.

Note that among these three policies DF appears as the most natural extension of the basic policy of Section 4.1. As, with DF, idle time is allocated only if all customers that have been confirmed for service so far have been served already, decision makers do not have to worry about possibly having to reject a late requesting customer due to having spent too much idle time at an earlier point in time. Moreover, given that the final destination of a vehicle usually is a depot, a transportation hub or a warehouse, each of which is typically located rather far away from potential customer locations, it may be concluded that the DF policy will always be at least as good as the basic policy. Nevertheless, Section 6 shows that DF is

the best choice for only a subset of the possible application scenarios.

5 Posteriori Evaluation Approach

The class of problems described in Section 3 falls under the broad umbrella of dynamic vehicle routing problems. Due to the uncertainty about customer requests we aim at policies that maximize only the expected number of confirmed requests over the considered time period (cf. equation 6).

At the end of such a time period, however, we are able to determine what the ideal solution would have been, if we had known in advance the points in time at which the dynamic customer requests occurred. In Section 6.2 we compare ideal a posteriori solutions to solutions resulting from application of the policies introduced in Section 4. The remainder of the present section describes the problem to be solved in order to determine the quality of ideal a posteriori solutions.

From an a posteriori point of view the points in time at which customer requests occur are known deterministically. Let $\mathcal{R} = \cup_t \mathcal{R}_t$ represent the set of customers that have issued a service request during the considered time horizon, and let t_i^r represent the point in time at which customer i has issued his request. Moreover, let \mathcal{E} represent the set of road links (i, j) connecting any two locations i and j of set $\mathcal{R} \cup \mathcal{C}_0 \cup \{1, N\}$. The problem to be solved for determination of the ideal posteriori solution may then be defined by the following Equations 7a-7i:

$$\max \sum_{i \in \mathcal{R}} x_i^c \quad (7a)$$

$$s.t. \quad x_i^c = 1 \quad \forall i \in \mathcal{C}_0 \cup \{1, N\}, \quad (7b)$$

$$\sum_{(i,j) \in \phi^+(i)} x_{ij}^p = x_i^c \quad \forall i \in \mathcal{R} \cup \mathcal{C}_0 \cup \{1, N\}, \quad (7c)$$

$$\sum_{(k,i) \in \phi^-(i)} x_{ki}^p = x_i^c \quad \forall i \in \mathcal{R} \cup \mathcal{C}_0 \cup \{1, N\}, \quad (7d)$$

$$\sum_{(k,i) \in \phi^-(i)} y_{ki} + \sum_{(k,i) \in \phi^-(i)} d_{ki} x_{ki}^p \leq \sum_{(i,j) \in \phi^+(i)} y_{ij} \quad \forall i \in \mathcal{R} \cup \mathcal{C}_0 \cup \{N\}, \quad (7e)$$

$$\sum_{(i,j) \in \phi^+(i)} t_i^r x_{ij}^p + \sum_{(k,i) \in \phi^-(i)} d_{ki} x_{ki}^p \leq \sum_{(i,j) \in \phi^+(i)} y_{ij} \quad \forall i \in \mathcal{R} \cup \mathcal{C}_0 \cup \{1, N\}, \quad (7f)$$

$$T x_{ij}^p \geq y_{ij} \quad \forall (i, j) \in \mathcal{E}, \quad (7g)$$

$$y_{ij} \in \mathbb{N}_0 \quad \forall (i, j) \in \mathcal{E}, \quad (7h)$$

$$x_{ij}^p, x_i^c \in \{0, 1\} \quad \forall (i, j) \in \mathcal{E}, \forall i \in \mathcal{R} \cup \mathcal{C}_0 \cup \{1, N\} \quad (7i)$$

We aim at maximizing the number of confirmed customer requests over the entire time horizon (7a) while visiting start location, final destination as well as each of the advance customers that are already confirmed at $t = 0$ (7b). Confirmation decisions x_i^c and routing decisions x_{ij}^p are defined corresponding to their time-dependent counterparts in Section 3.2. Equations 7c and 7d make sure that start and final destination as well as each customer location with a confirmed request is both approached and left exactly once. Each of the variables y_{ij} is equal to zero if link (i, j) is not part of the solution, i.e., $y_{ij} = 0$ if $x_{ij}^p = 0$. In case of $x_{ij}^p = 1$, variable y_{ij} represents the point in time at which the vehicle arrives at location i .

Equation 7e makes sure that the vehicle never arrives at location i before it's predecessor k (according to the route) has been visited and the distance d_{ki} has been travelled. As a consequence we have $y_{ki} \leq y_{ij}$ with j representing the successor of i . Additionally, Equation 7f makes sure that the vehicle is never allowed to leave the preceding location k for moving on to i before i has actually issued a service request. By requiring $T x_{ij}^p \geq y_{ij}$ (7g) the model is guaranteed to arrive at the final destination by time $t = T$ with all confirmed requests being served.

The given model may be classified as representing a deterministic travelling salesman problem with time windows and selective service. TSPs with selective service are sometimes also referred to as "Orienteering Problems" in the literature (e.g., Kantor and Rosenwein, 1992). Note that the above model differs from more traditional models for the orienteering problems in terms of the claim made by Equation 7f. The above model represents a novel approach to orienteering problems that has, however, been inspired by Eijl's model for the "Delivery Man Problem with Time Windows" (Eijl, 1995).

Solving the model with a given set \mathcal{R} of requesting customers and with their request times t_i^r , unveils what the ideal decisions for the corresponding time period would have been as well as how many confirmations of dynamic requests one could have made at most. We will use these figures in Section 6.2 for comparing the idle time policies proposed in Section 4 with posteriori optimal solutions.

6 Empirical Evaluation

We carried out a large number of computational experiments in order to systematically evaluate the performance of the policies proposed in Section 4.2, and in order to be able to derive sound criteria for applying these policies within routing information systems. The experiments included systematic variation of all relevant problem characteristics. In particular we varied the following characteristics:

- **service region:** A number of different service regions were considered, each of which we derived from well-known standard vehicle routing problem instances (Solomon, 1987; TSPLIB, 2014). The regions differ with respect to the geographical arrangement of a total of 50 customer locations. Both the case of customer locations being evenly distributed over the service region and the case of customer locations being grouped into several clusters within the region was considered. Additionally we considered intermediate cases between the two.
- **time horizon:** As a consequence of the number of customer locations being fixed at 50, we varied the number T of time units available for each service region between a minimum value (hardly allowing for confirmations, even if $|\mathcal{R}| = 1$) and a maximum value (hardly requiring rejections even if $|\mathcal{R}| = |\mathcal{S}^d|$).
- **advance customers:** For each choice of service region and time horizon we did experiments with a variety of ratios $\frac{|\mathcal{S}^a|}{|\mathcal{S}|}$ of advance customers as they typically occur in practice. The customer locations representing the advance customers were selected randomly for each experiment.
- **dynamic customers:** For each experiment the ratio of dynamic customers is $1 - \frac{|\mathcal{S}^a|}{|\mathcal{S}|}$. Each dynamic customer is assumed to issue at most one single service request within the given time horizon. The points in time at which requests occur are determined by the interarrival times of a Poisson process. The actual customer requesting for service at request time t is selected randomly from the set of the dynamic customers that have not yet issued a service request.

With each problem instance, consisting of service region, time horizon, a ratio of advance customers as well as a Poisson process determining arrival times of dynamic service requests, we considered a total of 250 different sample realizations. Note that a sample realization determines (1) the selection of advance customers, (2) the points in time of service requests and (3) the selection of dynamic requests. For comparison we averaged the performance of each of WF, DF and DW over the 250 sample realizations. For maximizing the number of confirmations and for solving open TSPs, we used a standard solver relying on branch-and-cut methods.

The idle time policies have been evaluated in two respects. Section 6.1 shows the key results of comparing the policies with each other. In particular 6.1 unveils the criteria for application of idle time policies in routing information systems. Section 6.2 evaluates the idle time policies with respect to ideal a posteriori solutions.

6.1 Relative Evaluation

The computational results indicate consistently that the geographical arrangement of customer locations within the service region is - independently of other problem characteristics - a key criterion for selection of the right idle time policy. Two example service regions are illustrated in Figure 1. The region shown in the left hand side subfigure of Figure 1 represents an extreme case of a geographical arrangement in the sense that customer locations are spread roughly evenly over the entire region. In contrast, the right hand side subfigure shows a service region consisting of five clusters of customer locations, with quite small intra-cluster distances between locations and relatively long inter-cluster distances. Note that service regions with clusters are quite common in real-world applications where customer locations may occur in a number of neighboring towns or in different districts of a city.

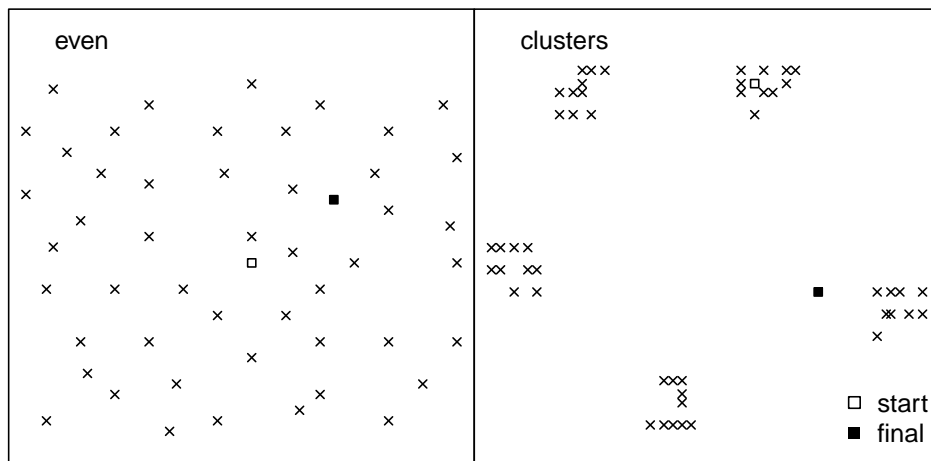


Figure 1. Two example service regions "even" and "clusters" featuring a total of $|\mathcal{S}| = 50$ customer locations each. Starting locations and final destinations are represented as squares.

The results of our experiments with a variety of service regions featuring different geographical arrangements clearly showed that DF outperforms both WF and DW in the absence of clear-cut customer clusters. Figure 2 provides an example comparison of the policies' performances with the two regions of Figure 1.

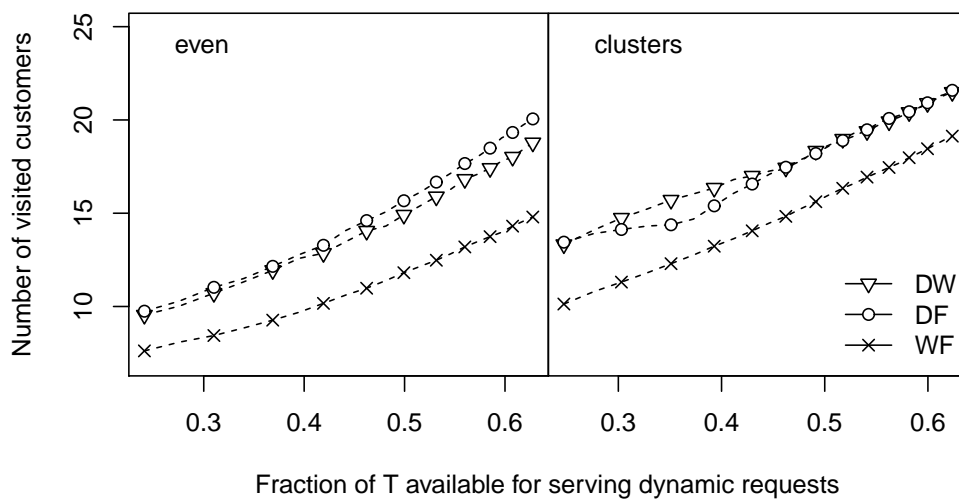


Figure 2. Average numbers of visited customers at different fractions $\frac{\tau_0}{T}$. Averages are shown for the idle time policies of Section 4.2 and for each of the two service regions of Figure 1.

Both of the two subfigures of Figure 2 display average numbers of visited customer locations over 250 sample realizations at different fractions $\frac{\tau_0}{T}$ of the entire time horizon being available for serving dynamic requests. Ratios of advance customers and of dynamic customers are equal for each sample realization.

The figure shows that, independent of the service region, each of the policies allows for more customer locations to be confirmed and visited as $\frac{\tau_0}{T}$ increases. For both service regions, WF performs significantly worse than DF and DW. In case of roughly even distribution of locations, DF performs better than DW no matter what $\frac{\tau_0}{T}$ is. In contrast, the performance of DW is greater or equal to the performance of DF if the service vehicle is routed in the service region with clusters.

The fact that for each service region the performances of DF and DW become similar as time for serving dynamic customers decreases, clearly is a result of DW and DF turning into very similar policies if τ_0 approaches 0. With evenly distributed locations, the advantage of DF over DW grows as $\frac{\tau_0}{T}$ increases, which indicates that DF is able to use the additional time efficiently, i.e., to use the additional time for visiting a relatively large number of late requesting dynamic customers that are close-by the planned route.

However, the advantage of DF vanishes in the presence of clusters. As shown in the right hand side subfigure of Figure 2, DF performs significantly worse than DW for $0.25 \leq \frac{\tau_0}{T} \leq 0.45$. Here DF suffers from the fact that it tends to move the vehicle quite rapidly into the next cluster of locations without having enough time for returning into previously visited clusters that might have new dynamic requests at a later point in time. In contrast, DW spends more time in earlier clusters and therefore tends to minimize the number of times it needs to return into previously visited clusters. As $\frac{\tau_0}{T}$ increases above 0.45 the (dis-)advantages of DW and DF seem to balance each other resulting into almost equal behavior of the two policies.

Figure 3 illustrates the advantage of DW in the clustered service region more clearly by displaying the performances of the three policies with respect to four alternative customer request behaviors.

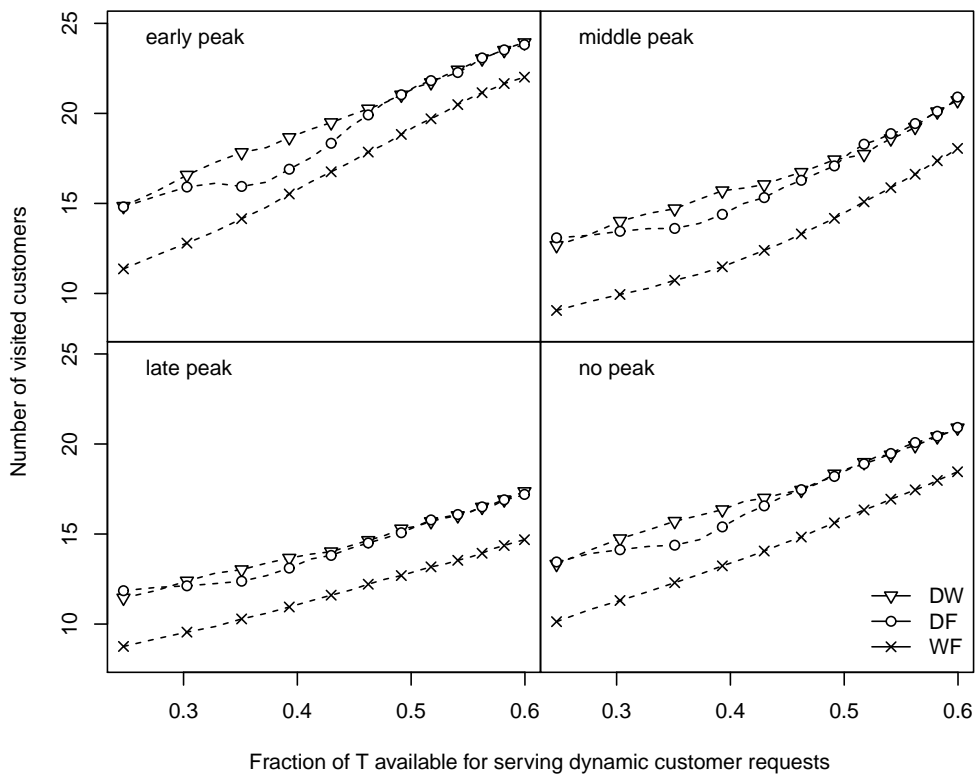


Figure 3. Average numbers of visited customers at different fractions $\frac{\tau_0}{T}$ with four different customer request behaviors in the clustered service region.

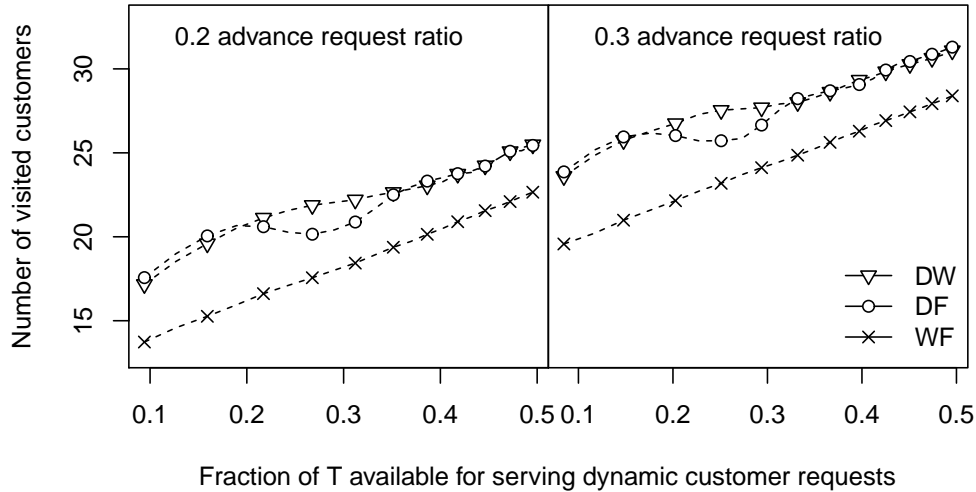


Figure 4. Average numbers of visited customers at different fractions $\frac{t_0}{T}$ with two different ratios of advance customer requests in the clustered service region.

The underlying problem instances differ from the ones of Figure 2 merely in terms of the points in time at which the dynamic customers issue their requests. We modified the instances by dividing time horizons into three intervals of equal size and by assigning customer request times to these intervals such that request peaks are simulated. The left top subfigure works with sample realizations where most of the requests occur in the first third of the time horizon, the right top subfigure has most requests in the second third of the horizon and the left bottom subfigure has most requests in the last third. The right bottom subfigure corresponds to the right hand side of Figure 2 and has an equal distribution of requests over the entire time horizon. Figure 3 illustrates that the advantage of DW over DF is the more significant the earlier most of the dynamic requests occur.

All insights illustrated up to this point have been independent of the ratio of advance requests. However, Figure 4 shows that the ratio of advance requests has a noticeable influence on the performance of the DF policy. The two subfigures of Figure 4 display results for problem instances that differ from the ones of Figure 2 merely in terms of the ratio of advance requests. Figure 2 is based on an advance request ratio of 0.1 which leads to a monotonically increasing average number of visited locations as the available time increases. In contrast, Figure 4 shows that increasing the advance request ratio leads to non-monotonic behavior of DF, i.e., in case of DF and a high ratio of advance customers, having more time available for visiting dynamic customers may lead to less dynamic customers being visited.

6.2 Posteriori Evaluation

The empirical results shown in Section 6.1 illustrate the main insights we gained from the large number of computational experiments we carried out for comparing the different idle time policies with each other. In order to be able to fully assess the value of these policies we additionally compared their performances with the qualities of the corresponding ideal a posteriori solutions. The ideal posteriori solutions have been determined by solving the model proposed in Section 5 with a standard solver relying on branch-and-cut methods.

As with the previous experiments, each of WF, DF and DW as well as the a posteriori model has been applied to a variety of problem instances featuring different service regions and different customer ratios. For each problem instance all four approaches were compared by averaging over 50 sample realizations at different ratios $\frac{t_0}{T}$.

As a main result we found that the quality of the best idle time policy tends to deviates from the a posteriori

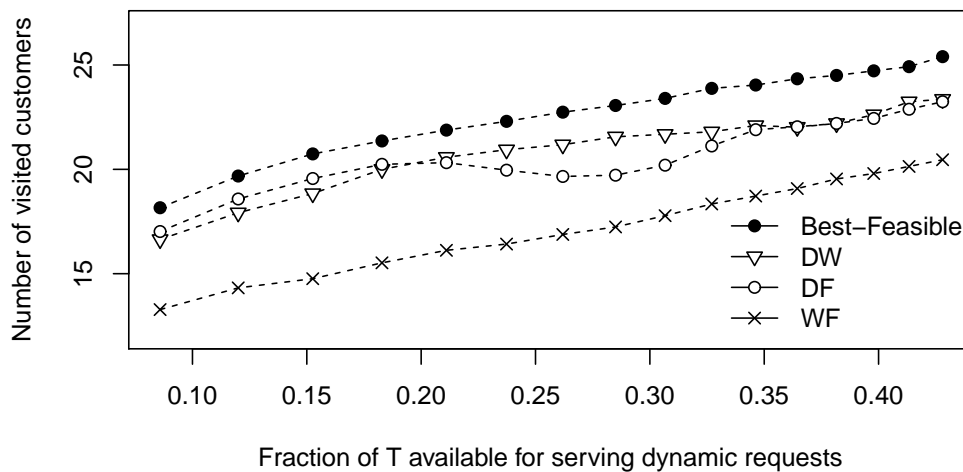


Figure 5. Average numbers of visited customers at different fractions $\frac{\tau_0}{T}$ for each of WF, DF and DW and for the corresponding ideal posteriori solutions (best feasible solution after 2000 seconds of branch-and-cut).

ideal solution to a larger extent, the higher the ratio $\frac{\tau_0}{T}$ is, i.e., the more time available for serving dynamic requests, the less efficient the idle time policy are. Depending on problem instance and time available, the average number of customers visited by the best idle time policy was at least between ten and twenty percent lower than the average number of customers visited by the ideal solution, respectively.

A representative set of results for a given problem instance with clustered customer locations and an advance customer ratio of 0.2 is displayed in Figure 5.

7 Conclusion

We empirically evaluated three alternative idle time policies (Wait First, Drive First, Distributed Wait) for real-time routing of a service vehicle. Our goal has been to provide both routing information system providers and supply chain operators with criteria for the use of idle time policies. The main characteristic of the class of routing problems considered is that part of the customer requests are uncertain, i.e., that part of the requests are not known in advance, but occur only while the service vehicle already is en-route. Both a dynamic model and a model allowing for determination of ideal a posteriori solutions to the routing problem were proposed and applied for the purpose of evaluating idle time policies. We carried out experiments with a variety of problem instances that differ in terms of service region, length of time horizon, ratio of advance requests as well as request behavior of dynamic customers.

The empirical results led to a number of conclusions that may serve as criteria for the use of idle time policies for the considered problem class: In particular, we have shown that the Wait First policy is always outperformed by the other policies. Moreover, the Drive First policy turned out to be more efficient than the other policies if no clear-cut clusters can be identified in the service region. However, in the presence of clusters the Distributed Wait policy is more efficient, especially in case of many early dynamic requests. Drive First shows significant disadvantages with clustered customer locations, in case of a high ratio of advance requests. Comparison with the quality of a posteriori solutions showed that following the best idle time policy leads to between ten and twenty percent less visited customers than theoretically possible. Many avenues for future work exist. More sophisticated idle time policies may lead to better results and in order to better match the situation of many real-world supply chains the present study could be extended by considering the value of idle time policies in the presence of additional sources of uncertainty such as random travel times.

References

- Bent, R. and P. van Hentenryck (2007). "Waiting and Relocation Strategies in Online Stochastic Vehicle Routing." In: *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*. Ed. by M. Veloso. Menlo Park, CA: AAAI Press, pp. 1816–1821.
- Bertsimas, D. and G. van Ryzin (1991). "A stochastic and dynamic vehicle routing problem in the Euclidean plane." *Operations Research* 39 (4), 601–615.
- (1993). "Stochastic and dynamic vehicle routing problem in the Euclidean plane with multiple capacitated vehicles." *Operations Research* 41 (1), 60–76.
- Branke, J. et al. (2005). "Waiting Strategies for Dynamic Vehicle Routing." *Transportation Science* 39 (3), 298–312.
- Chen, Z. and H. Xu (2006). "Dynamic column generation for dynamic vehicle routing with time windows." *Transportation Science* 40 (1), 74–88.
- Eijl, C. A. van (1995). *A polyhedral approach to the delivery man problem*. Tech. rep. Memorandum COSOR 95-19. Eindhoven University of Technology.
- Gendreau, M. et al. (1999). "Parallel tabu search for real-time vehicle routing and dispatching." *Transportation Science* 33 (4), 381–390.
- Ichoua, S. et al. (2000). "Diversion Issues in Real-Time Vehicle Dispatching." *Transportation Science* 34 (4), 426–438.
- (2006). "Exploiting knowledge about future demands for real-time vehicle dispatching." *Transportation Science* 40 (2), 211–225.
- Kantor, M. G. and M. B. Rosenwein (1992). "The Orienteering Problem with Time Windows." *The Journal of the Operational Research Society* 43 (6), 629–635.
- Larsen, A. et al. (2004). "The a priori dynamic traveling salesman problem with time windows." *Transportation Science* 38 (4), 459–472.
- Meisel, S. (2011). *Anticipatory Optimization for Dynamic Decision Making*. Vol. 51. Operations Research/Computer Science Interfaces Series. Springer New York.
- Mitrovic-Minic, S. and G. Laporte (2004). "Waiting strategies for the dynamic pickup and delivery problem with time windows." *Transportation Research Part B* 38 (7), 635–655.
- Mitrovic-Minic, S. et al. (2004). "Double-horizon based heuristics for the dynamic pickup and delivery problem with time windows." *Transportation Research Part B* 38 (8), 669–685.
- Papastavrou, J. (1996). "A stochastic and dynamic routing policy using branching processes with state dependent immigration." *European Journal of Operational Research* 95 (1), 167–177.
- Pillac, V. et al. (2013). "A review of dynamic vehicle routing problems." *European Journal of Operational Research* 225 (1), 1–11.
- Powell, W. (2011). *Approximate Dynamic Programming*. Hoboken, NJ: John Wiley & Sons.
- Solomon, M. M. (1987). "Algorithms for the Vehicle Routing and Scheduling Problems with Time Window Constraints." *Operations Research* 35 (2), 254–265.
- Swihart, M. and J. Papastavrou (1999). "A stochastic and dynamic model for the single-vehicle pick-up and delivery problem." *European Journal of Operational Research* 114 (3), 447–464.
- Thomas, B. (2007). "Waiting Strategies for Anticipating Service Requests from Known Customer Locations." *Transportation Science* 41 (3), 319–331.
- TSPLIB (2014). *Library of sample instances for the TSP*. University of Heidelberg, Germany. URL: <http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95>.