Available online at www.sciencedirect.com

ScienceDirect

Transportation Research Procedia 10 (2015) 364 – 373

Transportation Research Procedia

www.elsevier.com/locate/procedia

18th Euro Working Group on Transportation, EWGT 2015, 14-16 July 2015, Delft, The Netherlands

# Short-term strategies for stochastic inventory routing in bike sharing systems

Jan Brinkmann*, Marlin W. Ulmer, Dirk C. Mattfeld

*Decision Support Group, Technische Universität Braunschweig, Mühlenpfordtstr. 23, 38106 Braunschweig, Germany*

**Abstract**

Bike sharing systems (BSS) provide individual and eco-friendly urban mobility and are implemented in a growing number of cities. In BSS, customers can rent and return bikes spontaneously at stations and at every time of the day. To allow a reliable usage, system operators have to enable a sufficient number of bikes and empty bike racks at each station. Therefore, system operators use a set of vehicles to relocates bikes between stations. The according routing can be derived solving an inventory routing problem (IRP). For planning, operators can draw on expected customer trips generally following specific daytime patterns. Nevertheless, a significant amount of rentals and returns occur unpredictably and spontaneously forcing immediate adaptions of the routes.

In this paper, we define the stochastic IRP for BSS and present a short-term relocation strategy (STR). A STR defines priority stations regarding their urgency that have to be rebalanced. In a real world case study, we compare STR to a long-term relocation strategy (LTR) using given target fill levels. STR outperforms LTR significantly leading to suitable service levels.



Peer-review under responsibility of Delft University of Technology

*Keywords:* inventory routing; vehicle routing; stochastic dynamic programming; shared mobility; bike sharing

## 1. Introduction

Due to urbanization, the volumes of individual motorized traffic increase in cities worldwide (McCarthy and Knox, 2005). The results are traffic jams and environmental pollution. To reduce the number of motorized vehicles, city authorities draw on public transport and shared mobility systems. In particular, public bike sharing systems (BSS) are implemented to allow individual and eco-friendly transportation (Büttner et al., 2011). In BSS, a set of stations is distributed in a city. Each station contains a number of bike racks and bikes. At every station, users can rent and return bikes spontaneously. Typically, rental and return stations are not identical, i.e., one-way trips are usual. Due to spatio-temporal variation of user requests, stations tend to run full or out of bikes in the course of the day (Vogel et al., 2011). Further, due to the spontaneous and uncertain rental and return behavior, arbitrary stations might run empty or full unpredictably. Empty and full stations lead to failing user requests. At empty stations, requests for

* Corresponding author. Tel.: +49 531 391-3212 ; fax: +49 531 391-8144.
*E-mail address:* j.brinkmann@tu-braunschweig.de



Peer-review under responsibility of Delft University of Technology
doi:10.1016/j.trpro.2015.09.086

bikes cannot be satisfied while at full stations, requests for empty racks cannot be satisfied. A failed request leads to customer dissatisfaction and even a rejection of the whole BSS concept. So, system operators have to ensure a sufficient number of bikes and empty bike racks at each station for each point in time to reliably satisfy user requests. To achieve a sufficient service level, i.e., a high percentage of fulfilled customer requests, the system operator relocates bikes between stations with a fleet of vehicles. So, the system operator has to decide about the stations to serve and the number of bikes to relocate.

The presented problem can be formulated as an inventory routing problem with unknown due dates. Given a set of stations and inventory fill levels, the dispatcher routes capacitated vehicles between the stations. Inventory decisions are made about increasing or decreasing fill levels at each station to avoid customer dissatisfaction. A dissatisfaction occurs if a customer requires an inventory item (bike) or an empty inventory space (bike rack). Customer requests induce due dates for every station. A due date is the latest time a station's fill level can be adapted in order to fulfill requests. A due date is violated if the inventory is empty or full and an according request occurs. Since requests are uncertain, due dates are unknown as well. The objective is to minimize the number of due date violations over the planning horizon.

For planning, system operators can draw on current fill levels in stations and on expected future trips offered by data analysis (Borgnat et al. 2011, Vogel et al. 2011). A trip consists of one rental request and one return request. Furthermore, external information systems provide target fill levels as input for IRP (Schuijbroek et al. 2013, Vogel et al. 2014). These target fill levels are anticipatory, since information about expected future requests are considered. E.g., high target fill levels are determined if a high number of bike requests is expected. Target fill levels can be realized by transport vehicles (Raviv et al. 2013, Kloimüllner et al. 2014, Brinkmann et al. 2015).

In this article, we study the trade-off between the number of served stations and the ration of relocation operations and served stations. Thus, we introduce a short-term relocation strategy (STR). This strategy selects stations regarding their urgency and their immediate violation risk without making use of expected future requests. We compare the short-term relocation strategy with the long-term relocation strategy (LTR) by Brinkmann et al. (2015), which realizes anticipatory target fill levels given by Vogel et al. (2014). There are notable differences between STR and LTR: Firstly, STR neglects information about future requests while LTR rests upon expected requests offered by data analysis. Secondly, STR serves many stations and relocates a small number of bikes at each served station while LTR serves a small number of stations and relocates a large number of bikes at each served station. In a real world case study, we show that for the given problem, STR allows a substantially higher service level reducing the number of due date violations up to 72.52%.

This article is structured as follows: In Section 2, we refer to literature on both vehicle routing and bike sharing systems. A definition of the inventory routing problem in bike sharing systems is given in Section 3. The short-term and the long-term relocation strategies are introduced in Section 4. In Section 5, real world case studies are presented. This work concludes and gives an outlook of possible future work in Section 6.

## 2. Literature Review

The literature on vehicle routing and inventory management problems is vast. In vehicle routing problems (VRP), a fleet of vehicles moves between customer locations. The objective typically is to visit each location once within minimum time (Laporte, 1992), or to visit as many locations as possible within a limited time horizon (Ulmer et al., 2015). A multi-periodic VRP with due dates is described by Archetti et al. (2015). Here, a due date is the latest time a customer has to be served. All due dates are known in advance. Inventory management problems try to satisfy a given consumption while minimizing variable inventory and fixed order costs. A combination is introduced by Dror et al. (1985) and is called inventory routing problem (IRP). In an IRP, capacited vehicles serve customers. Each customer has a certain consumption of a commodity. The challenge is to provide the commodity in order to satisfy each user's consumption. Here, costs regarding transportation, inventory, and/or unsatisfied requests have to be minimized. A large overview of different IRPs is provided by Coelho et al. (2014).

The literature on BSS focuses on decision support for system operators. A system has to be installed and maintained to guarantee its functionality at all times in a cost efficient way (Benchimol et al., 2011). Literature on BSS can be divided into two groups. The first group aims on data analysis and optimization for determining optimal fill levels. Data analysis offering insights into general BSS behavior has been done by Borgnat et al. (2011) and Vogel et al.

(2011). Findings are patterns in the spatio-temporal distributions of trips. These patterns can be used to generate suitable fill levels for each station per hour. Schuijbroek et al. (2013) use Markov chains to define fill levels. A resource allocation problem minimizing costs regarding relocations and unsatisfied requests is introduced by Vogel et al. (2014).

The second group aims on IRP in the context of BSS. Relocation operations are typically done with respect to given target fill levels (Brinkmann et al., 2015). Nevertheless, IRP for BSS differs in a number of features. Chemla et al. (2013) uses a single vehicle to rebalance a system while minimizing the tour length. A problem setting including a fleet of vehicles and a limited time horizon is introduced by Raviv et al. (2013), who minimize the tour length and the difference of target fill levels and realized fill levels, i.e., the gap. More complex multi-criteria objective functions are used by Di Gaspero et al. (2013) and Rainer-Harbach et al. (2013). Here, gaps, tour length, and relocation operations are minimized. Deterministic user requests are considered by Kloimüllner et al. (2014), who minimizes the number of unsatisfied requests. For a detailed literature review on IRP in BSS, we refer to Brinkmann et al. (2015).

## 3. A Stochastic Inventory Routing Problem for Bike Sharing Systems

In this section, we describe and define the stochastic inventory routing problem for bike sharing systems using a Markov Decision Process. For better understanding, an example is given.

### 3.1. Problem Description

Consider a BSS consisting of stations, bikes, capacited vehicles and a depot. Each station has an initial fill level and a limited number of bike racks. Regarding the working day, all vehicles start and end their tours at a depot. Vehicles start empty at the depot in the first point in time and have to return empty at the end of a given time horizon. Stations may be served multiple times by one vehicle and different vehicles. We consider a service time per relocated bike, i.e., a time needed for moving one bike from a vehicle into a station or from a station onto a vehicle. Over time, uncertain trips occur. A trip consists of a rental and a return request. If a requests fails due to an empty or full station, the according due date is violated and the request has to be repaired. I.e., the user is sent to the nearest station where the associated request can be satisfied. The objective is to minimize the number of due date violations.

### 3.2. Problem Definition

Let $N = \{n_0, n_1, \ldots, n_{\max}\}$ be a set of nodes representing the stations and a depot $n_0$. Every station has a capacity $r\colon N \to \mathbb{N}_0$, i.e., a limited number of bike racks. Further, a set of edges $E = \{e_{ij}\colon n_i, n_j \in N\}$ connect stations with according travel times $d\colon E \to \mathbb{R}^+$. A set of vehicles $V = \{v_1, v_2, \ldots, v_{\max}\}$ is given. Vehicles are capacitated with $c\colon V \to \mathbb{N}$. The time horizon is defined as $T = \{0, \ldots, t_{\max}\}$. Customer requests $\sigma\colon N \times T \to \mathbb{Z}$ are stochastic and can occur at every station in every point in time. If $\sigma(n,t) < 0$, rental requests occur at station $n$ in time $t$ while $\sigma(n,t) > 0$ indicates return requests. Since each bike rented has to be returned within the closed system, $\sum_{n\in N} \sum_{t\in T} \sigma(n,t) = 0$ holds for all instances.

The problem described is stochastic and dynamic as defined by Kall and Wallace (1994). It is dynamic because decisions can be adapted over the planning horizon. The problem is stochastic because trips follow a certain distribution and are not known in advance. A stochastic dynamic decision problem can be defined using a Markov Decision Process (MDP) based on Bellman (1957), depicted in Formula (1).

$$S_k \quad \xrightarrow{x} \quad S_k^x \quad \xrightarrow{\omega} \quad S_{k+1} \tag{1}$$

In a MDP, decisions are made for a number of decision points $k \in T$. In every decision point $k$, a decision state $s_k \in S = \{s_0, s_1, \ldots, s_{max}\}$ and a set of possible decisions $X(s_k) = \{x_0, x_1, \ldots, x_{max}\}$ are given. In every decision state, a policy $\pi\colon S \to X(s_k)$ offers a decision. Every decision $x$ leads to a deterministic post-decision state $s_k^x$. A stochastic transition $\omega\colon S \times X(s_k) \to S$ leads to the next decision state $s_{k+1}$. Decision states $s_k$, decisions $x$, and transitions $\omega$ can provide rewards or penalties $p\colon S \times X(s_k) \times \omega \to \mathbb{Z}$.

In the following, we use the MDP to define the IRP. For the given problem, a decision point $k$ occurs, when a vehicle arrives or stays at a station. A decision state is distinctly identified by a point in time $t \in T$, station fill levels

$f\colon N \rightarrow \mathbb{N}_0$, vehicle's next positions, corresponding arrival times, and vehicle loads $l\colon V \rightarrow \mathbb{N}_0$, i.e., the numbers of bikes loaded by vehicles. In each decision state $s_k$, a decision $x \in X(s_k)$ has to be made. Decisions are on the one hand about the inventory, i.e., the number of bikes to pick-up or to deliver – and on the other hand about the routing, i.e., the next station to serve. These decisions include idling at the current station. The deterministic post-decision state $s_k^x$ contains time $t$, station fill levels $f$, vehicle loads $l$, vehicle positions and corresponding arrival times. The transition to the next decision state $s_{k+1}$ realizes inventory and routing decisions until the next vehicle arrives at a station. This includes both travel times and the service time. Further, $\omega$ reveals requests $\sigma$ in this time span and therefore changes station's fill levels. Let $\upsilon$ be the number of due date violations within the time span between decision points $s_k$ and $s_{k+1}$. Then the associated penalty $p(s_k, x, \omega) = \upsilon$ occurs and the requests are served from alternative stations.

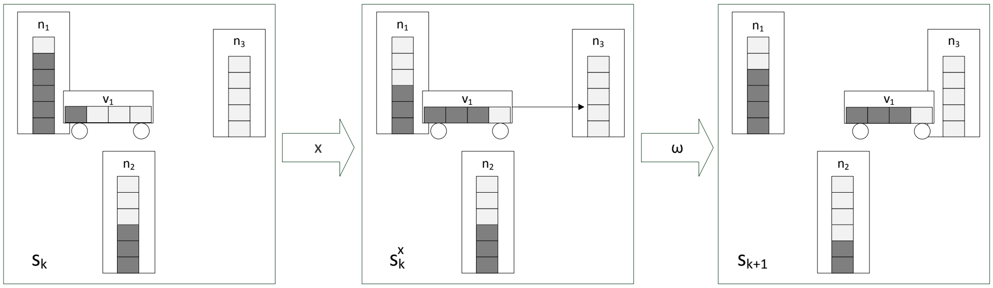


Fig. 1. Exemplary State, Decision, Post-Decision State, Transition, and Resulting State

Figure 1 shows an exemplary MDP from state $s_k$ over post-decision state $s_k^x$ to state $s_{k+1}$. The system consists of three different stations and one vehicle. For stations and vehicle, light boxes represent empty bike racks, dark boxes represent bike racks filled with a bike. Stations $n_1$ and $n_2$ have a capacity of 6 racks, station $n_3$ of 5 racks. The vehicle capacity is 4. In state $s_k$, depicted on the left side of Figure 1, $n_1$ contains 5 bikes, $n_2$ contains 3 bikes, and $n_3$ is empty. The vehicle is located at $n_1$. Inventory decisions are about how many bikes to pick-up or deliver at the current station $n_1$. The routing decision is about where to travel next, or idling at the current station. The applied decision $x$ is to pick-up two bikes at $n_1$ and travel to $n_3$. The resulting post-decision state $s_k^x$ is depicted in the center of Figure 1. While traveling to $n_3$, the stochastic transition $\omega$ reveals a rental request $\sigma(n_3, 1) = -1$ at $n_3$ and a return request $\sigma(n_1, 1) = 1$ at $n_1$. Because station $n_3$'s fill level does not allow to serve rental requests, the customer selects $n_2$ as an alternative. This results in a penalty $p(s_k, x, \omega) = 1$ and the new state $s_{k+1}$, as shown on the right side of Figure 1.

Since requests are uncertain, due date violations cannot be minimized directly. Thus, a policy $\pi$ aims on minimizing the expected number of due date violations. Let $s_0$ be the initial decision state and let $\Pi = \{\pi_0, \ldots, \pi_{max}\}$ be the set of all policies $\pi_i$. As shown in Formula (2), the objective is to identify an optimal policy $\pi^*$ that leads to the minimum of expected penalties, i.e., due date violations.

$$\pi^* = \underset{\pi \in \Pi}{\arg\min}\, \mathbb{E}\left[\sum_{k=0}^{T} p(s_k, \pi(s_k), \omega)\middle| s_0\right] \tag{2}$$

## 4. Strategies

Since it is computationally intractable to identify an optimal policy, we consider heuristic policies realize suitable fill levels to minimize the number of expected due date violations.

We define the short-term relocation strategy and recall the long-term relocation strategy introduced by Brinkmann et al. (2015). In both strategies and in each decision point, decisions regarding inventory and routing have to be made. The inventory decision affects the stations where the vehicle just arrives and depicts the number of bikes to relocate. Here, the strategies aim for realization of suitable fill levels. The assumption is that a suitable fill level fulfills due dates. For routing, each station is assigned a score. The score indicates whether a station is balanced or not, i.e., if its fill level is suitable in the sense of the strategy or not, and which station will be served next.

Since we consider travel time and service time as well, a trade-off between the number of served stations and the ratio of relocated bikes and served stations is given. I.e., on the one hand side, a strategy may serve as many stations as possible but ends in a small ratio of relocation operations and served stations – or on the other hand side, it conducts as many relocation operations at each served station and ends in a small number of served stations. The short-term relocation strategy aims on providing safety buffers to avoid due date violation in the near future. In order to save time, STR relocates a small number of bikes per station. The saved time is spent on serving more stations. Thus, at many stations, safety buffers will be realized. The long-term relocation strategy tries to realize given target fill levels anticipate future requests. This will lead to a large number of relocated bikes per served station while the number of served stations will be small. The assumption is, that if a target fill level at a certain station is provided, this station does not have to be served again in the future.

### 4.1. Short-term Relocation

The short-term relocation strategy (STR) follows a modified *first come, first serve*-approach. Stations with a certain urgency and therefore a high risk of immediate violations are prioritized. STR explicitly acts myopic, neglecting all information about expected future requests. To satisfy as many requests as possible, safety buffers $b\colon N \to \mathbb{N}$ are implemented. Safety buffers indicate the minimum number of bikes and empty racks that is desired at the given station for all points in time. If a station's fill level fulfills both safety buffers for bikes and for empty racks, the station is called balanced. If a safety buffer is violated, the station is called unbalanced. At unbalanced stations, we expect a high probability of failed requests in the near future. STR aims on realizing safety buffers as quick as possible. Therefore, the nearest unbalanced station is considered to be served by a vehicle. The number of relocated bikes is neglected.

When a decision point occurs, i.e., a vehicle $v$ arrives or stays at station $n_i$, an inventory decision has to be made. At first, the shortage of bikes or empty bike racks is represented by $\delta^{str}$ and determined according to Formula (3). It serves as preliminary inventory decision and depends on $n_i$'s fill level $f(n_i)$, safety buffers $b(n_i)$, and capacity $r(n_i)$.

$$\delta^{str} = \begin{cases} b(n_i) - f(n_i) : f(n_i) < b(n_i) \\ r(n_i) - f(n_i) - b(n_i) : r(n_i) - f(n_i) < b(n_i) \\ 0 : \text{else} \end{cases} \tag{3}$$

If $n_i$'s fill level is smaller than the safety buffer, i.e., $f(n_i) < b(n_i)$, $\delta^{str} > 0$ holds and indicates the shortage of bikes. Thus, bikes have to be delivered. If $n_i$'s number of empty bike racks is smaller than the safety buffer, i.e., $r(n_i) - f(n_i) < b(n_i)$, $\delta^{str} < 0$ holds and indicates the shortage of empty bike racks. Thus, bikes have to be picked-up. If $n_i$' fill level does not violated any safety buffer, $\delta^{str} = 0$ holds and indicates no need of relocation operations.

The realized number of relocated bikes $\iota^{str}$, i.e., the final inventory decision, has to be made respecting the vehicle load $l(v)$ and capacity $c(v)$ following Formula (4).

$$\iota^{str} = \begin{cases} \min\{\delta^{str}, l(v)\} : \delta^{str} > 0 \\ \max\{\delta^{str}, l(v) - c(v)\} : \delta^{str} < 0 \\ 0 : \text{else} \end{cases} \tag{4}$$

If $n_i$ has a shortage of bikes, i.e., $\delta^{str} > 0$, the number of delivered bikes is the minimum of $\delta^{str}$ and $l(v)$. I.e., the number of delivered bikes is either the shortage of bikes or the number of bikes loaded by the vehicle. If $n_i$ has a shortage of empty bike racks, i.e., $\delta^{str} < 0$, the number of picked-up bikes is the maximum of $\delta^{str}$ and $l(v) - c(v)$. I.e., the number of picked-up bikes is either the shortage of empty bike racks or the number of empty bike racks on the vehicle. These conditions explicitly avoid delivery operations larger than the number of bikes loaded by the vehicle, and pick-up operations larger than the vehicle's free capacity.

To make the routing decision, for each station $n_j \in N$ a score $\rho_j^{str}$ is determined. A score $\rho_j^{str}$ depends on whether station $n_j$ is balanced, whether the vehicle capacity $c(v)$ and number of loaded bikes $l(v)$ allow relocations, and on the inverse travel time between the vehicle's current station $n_i$ and $n_j$. $\rho_j^{str}$ is determined according to Formula (5).

$$\rho_j^{str} = \begin{cases} d(e_{ij})^{-1} : b(n_j) > f(n_j) \wedge l(v) > 0 \\ d(e_{ij})^{-1} : b(n_j) > r(n_j) - f(n_j) \wedge l(v) < c(v) \\ 0 : \text{else} \end{cases} \tag{5}$$

Thus, if vehicle $v$ can realize reposition operations, $\rho_j^{str}$ is the inverse travel time $d(e_{ij})^{-1}$, or 0 if vehicle $v$ cannot realize relocation operations.

The vehicle serves the nearest unbalanced station $n_p$, where relocations are possible. So, $\rho_p^{str}$ is the highest score as shown in Formula (6).

$$\rho_p^{str} \geq \rho_q^{str}, \forall n_q \in N \tag{6}$$

If $\rho_p^{str} = 0, \forall n_p \in N$, i.e., all stations are balanced or the only unbalanced station is the vehicle's current station, the vehicle does not travel to another station. This explicitly allows idling.

### 4.2. Long-term Relocation

The idea of the long-term relocation strategy (LTR) introduced by Brinkmann et al. (2015) is to realize target fill levels provided by external information systems (Schuijbroek et al. 2013, Vogel et al. 2014). These fill levels have been determined considering expected future requests. So, stations with a high number of expected future rental requests have high target fill levels. Stations with an expected high number of return requests have low target fill levels. This long-term relocation strategy is anticipatory because for decision making it takes information about future requests into account. Target fill levels are given by $\tau_h \colon N \to \mathbb{N}$, where $h$ is the hour of the day. We allow a tolerance $\mu \in \mathbb{N}$. If station $n$'s fill level lays in the target interval $[\tau_h(n) - \mu, \tau(n)_h + \mu]$, the station is called balanced. Else, it is unbalanced. For unbalanced stations, a positive gap indicates the distance between fill level and target interval. The idea is to minimize the gap over all stations. The station with the largest gap is considered to be served by a vehicle. The travel time is neglected because the assumption is, that a balanced station does not need to be served in the future again. When a station's gap is large, a large number of bikes need to be relocated.

The inventory decisions are made analogously to Formulas (3) and (4). When vehicle $v$ arrives or stays at station $n_i$, in Formula (7) the distance $\delta^{ltr}$ between $n_i$'s fill level and the target interval is determined. Thus, $\delta^{ltr}$ represents the shortage of bikes or empty bike racks.

$$\delta^{ltr} = \begin{cases} \tau_h(n_i) - \mu - f(n_i) & : \tau_h(n_i) - \mu > f(n_i) \\ \tau_h(n_i) + \mu - f(n_i) & : \tau_h(n_i) + \mu < f(n_i) \\ 0 & : \text{else} \end{cases} \tag{7}$$

$\delta^{ltr} > 0$ indicates a shortage of bikes, while $\delta^{ltr} < 0$ indicates a shortage of empty bike racks. As already shown for the short-term relocating strategy, the final inventory decision $\iota^{ltr}$ in Formula (8) additionally depends on the vehicle load $l(v)$ and capacity $c(v)$.

$$\iota^{ltr} = \begin{cases} \min\{\delta^{ltr}, l(v)\} & : \delta^{ltr} > 0 \\ \max\{\delta^{ltr}, l(v) - c(v)\} & : \delta^{ltr} < 0 \\ 0 & : \text{else} \end{cases} \tag{8}$$

So in the presence of a shortage of bikes, the number of delivered bikes is the minimum of $\delta^{ltr}$ and $l(v)$, i.e., either the shortage of bikes or the number of bikes loaded by the vehicle, or if $\delta^{ltr} < 0$, the maximum of $\delta^{ltr}$ and $l(v) - c(v)$, i.e., either the shortage of empty bike racks or the number of empty bike racks on the vehicle.

The routing decision is made according to scores $\rho_i^{ltr}, \forall n_i \in N$ as depicted in Formula (9). Scores $\rho_i^{ltr}$ base on the gap, i.e., the distance between target interval and fill level.

$$\rho_i^{ltr} = \begin{cases} \tau_h(n_i) - \mu - f(n_i) & : \tau_h(n_i) - \mu > f(n_i) \\ f(n_i) - (\tau_h(n_i) + \mu) & : \tau_h(n_i) + \mu < f(n_i) \\ 0 & : \text{else} \end{cases} \tag{9}$$

The vehicle serves the station $n_p$ with the largest gap according to Formula (10).

$$\rho_p^{ltr} \geq \rho_q^{ltr}, \forall n_q \in N \tag{10}$$

The special case $\rho_p^{ltr} = 0, \forall n_p \in N$ allows idling if all stations are balanced in the given decision point.

## 5. Case Studies

In this section, we describe the case study including instance generation and results. Regarding results we focus on the strategy's impact on relocation and routing.

### 5.1. Instances

The test instances are based on data analysis by Vogel et al. (2011) on a data set provided by Vienna's BSS "CityBike Wien" (Gewista Werbegesellschaft m.b.H., 2014). The data set contains approximately 750,000 single trips recorded on working days in the years 2008 and 2009. To that time, the BSS contained 59 stations and 627 bikes. Stations have a number of racks between 10 and 40. The number of bikes is equal to half of the number of bike racks over all stations. An artificial depot is placed at the most central located station.

The information system by Vogel et al. (2011) generates real-valued typical flows of bikes. I.e., for each pair of stations and hour, the probability for a trip is given. The real-valued flows are discretized by Poisson distribution. The outcomes are sets of in average 1,569 trips per day. As shown in Figure 2, a small morning peak of trips and a rush hour at the afternoon can be identified.

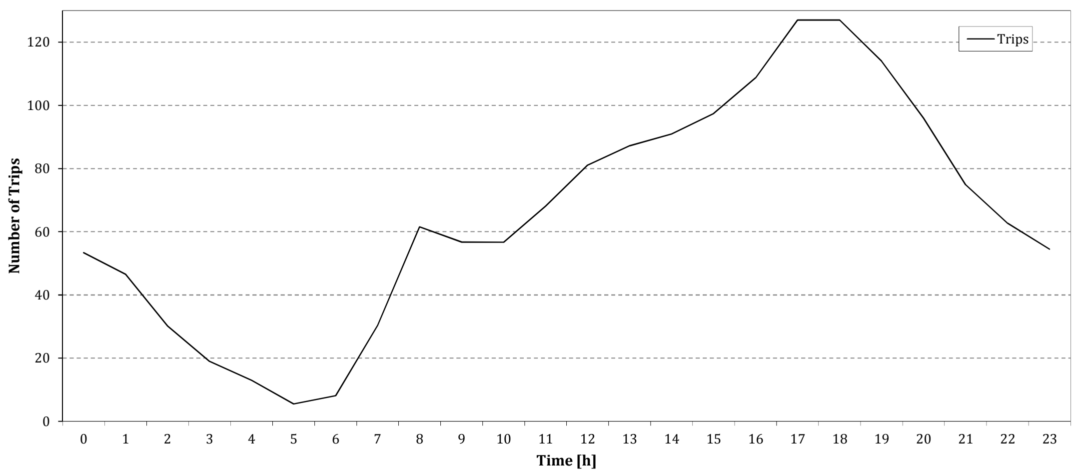


Fig. 2. Number of Trips per Hour for a Working Day

For the long-term relocation strategy, target intervals are given by Vogel et al. (2014). For both the short-term and the long-term relocation strategy initial fill levels are chosen randomly. I.e., for each bike a station is chosen where station's probabilities are equally distributed.

The system offers service 24 hours a day and 7 days a week. We are using one vehicle to rebalance the system. The constant speed is $15\frac{\text{km}}{\text{h}}$. The vehicle has a capacity of 10 bikes. Moving one bike from the vehicle into the station or from the station onto the vehicle takes a constant service time of 2 minutes. The distances between stations are assumed to be Euclidean. For the short-term relocation strategy, station's safety buffers are a fixed percentage $\beta$ of its capacity. I.e., station $n$'s safety buffers $b(n)$ depend on $\beta$ and the $n$'s capacity $r(n)$ as shown in Formula (11).

$$b(n) = \lceil \beta \cdot r(n) \rceil \tag{11}$$

For the long-term relocation strategy, the tolerance $\mu = 2$ is chosen. I.e., the target intervals have a size of 5.

For evaluating the strategies, 1,000 sets of trips are generated via Poisson distribution. Therefore, the simulation involves 1,000 non-concatenated working days for each strategy.

### 5.2. Results

We assume the average values to be a suitable approximation for expected values. E.g., the average number of due date violations is an approximation for the expected number of due date violations. Therefore, Table 1 shows the average results for the given strategy. It contains the average number of due date violations, the numbers of failed rentals and returns, the ratio of relocation operations and served stations. One relocation operation might be either moving one bike from the vehicle into the station or moving one bike from the station onto the vehicle. I.e., picking-up one bike at a station and delivering it to another station compromises two relocation operations. Additionally, the average ratio of relocation operations and served stations is shown. Simulations have been conducted for the long-term relocation strategy and short-term relocation strategy with different safety buffers. STR($\beta$) indicates relative safety buffers of $\beta$. The column "no relocations" shows results for a strategy without any relocations. These results serve as benchmark. The long-term strategy performs significantly better than the strategy without relocations. Here,

Table 1. Results

| | no relocations | LTR | STR(0.1) | STR(0.2) | STR(0.3) | STR(0.4) | STR(0.5) |
|---|---|---|---|---|---|---|---|
| due date violations | 131.252 | 98.630 | 53.916 | 39.987 | 36.068 | 41.847 | 64.997 |
| failed rentals | 36.129 | 30.745 | 7.095 | 4.393 | 6.357 | 9.816 | 17.565 |
| failed returns | 95.123 | 67.885 | 46.821 | 35.594 | 29.711 | 32.031 | 47.432 |
| relocation operations | – | 330.018 | 99.026 | 169.312 | 289.171 | 360.261 | 418.446 |
| served stations | – | 53.109 | 86.289 | 122.577 | 163.676 | 176.781 | 202.361 |
| ratio of relocation operations and served station | – | 6.214 | 1.148 | 1.381 | 1.767 | 2.038 | 2.068 |

the number of due date violations is decreased by 24.85%. However, the short-term strategy outperforms the long-term strategy by far. For the given scenario, a safety buffer of $\beta = 0.3$ offers a reduction of due date violations of 72.52%. For all strategies, the number of failed rentals is much smaller than the number of failed returns. This could indicate a surplus of bikes within the system. LTR leads to the smallest number of served stations and to the largest ratio of relocation operations and served stations. For STR the numbers of served stations and relocation operations, and ratio of relocation operations and served stations increase with the size of the safety buffers.

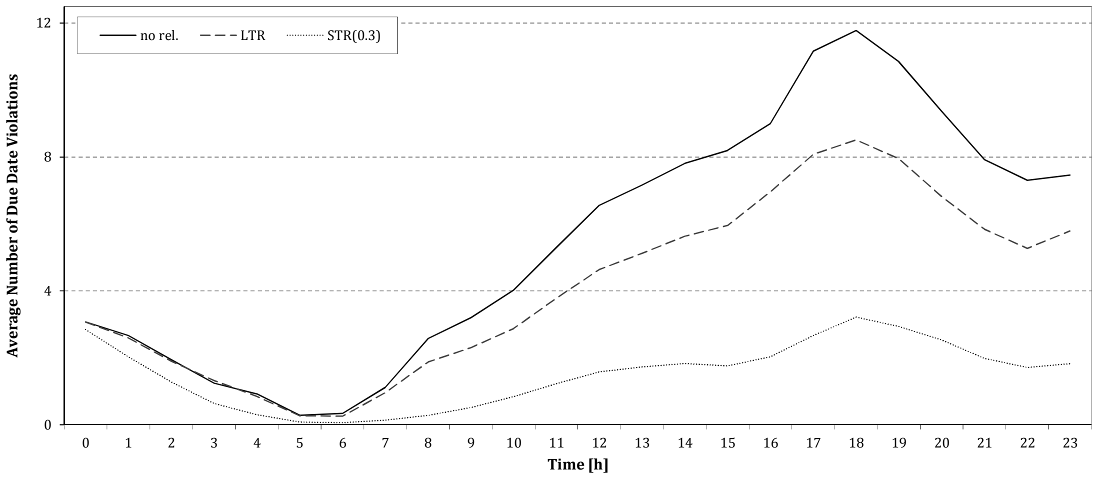


Fig. 3. Average Number of Due Date Violations per Hour for a Working Day

Figure 3 shows the average number of violated due dates for the system without relocations, the long term strategy and for the short-term strategy with a safety buffer of $\beta = 0.3$. For all strategies, the distribution of due date violations reflects the distribution of trips. As already seen in Table 1, the long-term strategy only provides small improvements. The short-term strategy is able to keep due date violations low even during high system usage in the afternoon.

To gain further insights into the strategies performances, we consider the service level. Let $S' = (s_0, \ldots, s_{max})$ be a sequence of decision states representing a simulation for a working day, and let $\pi(s_k) \in X(s_k)$ be the decisions offered by policy $\pi$ for the associated decision state $s_k$. Then the relative portion of satisfied requests for the given time horizon $T$ archived by policy $\pi$ is indicated by the service level $\lambda$ in Formula (12).

$$\lambda = 1 - \frac{\sum_{s_k \in S'} p(s_k, \pi(s_k), \omega)}{\sum_{n \in N} \sum_{t \in T} \left|\sigma(n,t)\right|} \tag{12}$$

The service level depicts the chance for satisfying a request within the time horizon. It offers an unbiased view on a strategies performance. Figure 4 shows the average service levels per hour. Again, the long-term relocation strategy can realize improvements in comparison to no relocations. After the morning hours, the service levels of these two strategies decrease constantly. I.e., the ratio of due date violations and trips increases and therefore the chance for satisfying a request decreases in the course of the day. For all hours, the short-term relocation strategy provides the highest service level. Expect for hours 0 to 2, the service level lies above 98%. This indicates a reliable service since the chance for satisfying a request does not change even in peak hours.

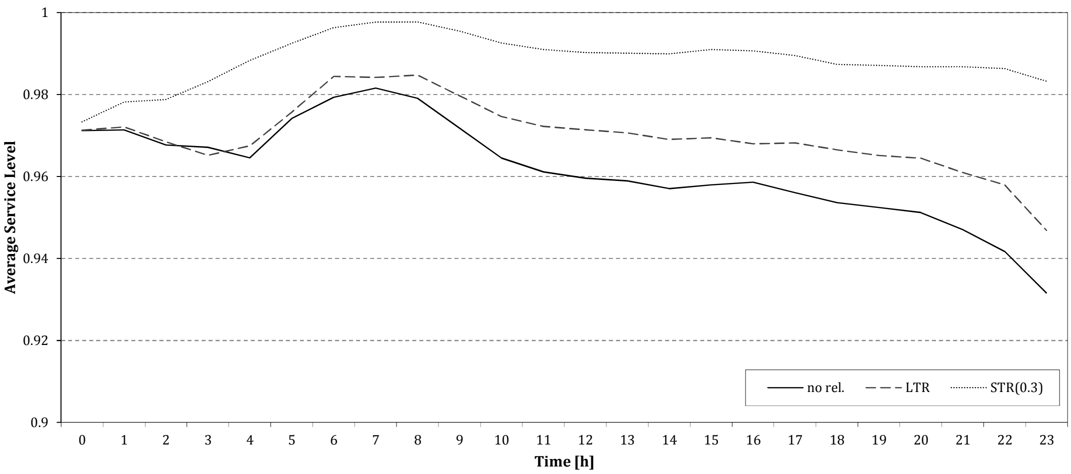


Fig. 4. Average Service Level per Hour for a Working Day

## 6. Conclusion

We have introduced a stochastic inventory routing problem for rebalancing bike sharing systems. This problem is formulated as an inventory routing problem with unknown due dates. A due date is the latest time a station has to be served by a transport vehicle relocating bikes in order to serve requests. Since customer requests are uncertain, the due dates are uncertain as well. The objective is to minimize the number of due date violations.

For solving these problem, a number of optimization algorithms aim on realizing given target fill levels at each station on the basis of expected trips. Therefore, information systems generating target fill levels have to be used. Long-term relocation strategies (LTR) use target fill levels for making decisions regarding inventory and routing. We also introduce a short-term relocation strategy (STR) considering a station's urgency by using safety buffers in order to decide whether a station has to be served or not. A safety buffer serves as the minimum number of bikes and free bike racks. If a buffer is violated, the station will be served by a vehicle to relocate bikes. STR and LTR depict a trade-off between the number of served stations and the ratio of relocation operations and served stations.

Our test results on a data set of Vienna's BSS "CityBike Wien" point out that both the short-term and the the long-term relocation strategy can decrease the number of due date violations significantly. If suitable safety buffer are chosen, STR outperforms LTR by far. Even during peak hours, the STR provides a small number of due date violations. Thus, we conclude that a station's urgency cannot be neglected.

Future research should concentrate on a comparison between short-term and long-term relocation strategies in much larger systems. Also an evaluation by routing a fleet of vehicles might be possible. A combination of long-term and short-term strategies might be essential if large system cannot be maintained well by a small number of vehicles. Further, time-depended safety buffers could be useful if requests at a station differ significantly in the course of the day.

## Acknowledgements

The authors thank Patrick Vogel for the provision of input data, and Justin C. Goodson for his helpful suggestions and comments in preparation of this work. We also thank Gewista Werbegesellschaft mbH for providing the required data from their project "CityBike Wien".

## References

Archetti, C., Jabali, O., Speranza, M.G., 2015. Multi-period Vehicle Routing Problem with Due dates. Computers & Operations Research 61, 122–134.

Bellman, R., 1957. A Markovian Decision Process. Technical Report. DTIC Document.

Benchimol, M., Benchimol, P., Chappert, B., De La Taille, A., Laroche, F., Meunier, F., Robinet, L., et al., 2011. Balancing the stations of a self-service bike hire system. RAIRO-Operations Research 45, 37–61.

Borgnat, P., Abry, P., Flandrin, P., Robardet, C., Rouquier, J.B., Fleury, E., 2011. Shared bicycles in a city: A signal processing and data analysis perspective. Advances in Complex Systems 14, 415–438.

Brinkmann, J., Ulmer, M.W., Mattfeld, D.C., 2015. Inventory Routing for Bikes Sharing Systems. Working Paper (2015-01-12).

Büttner, J., Mlasowsky, H., Birkholz, T., et al., 2011. Optimising Bike Sharing in European Cities - A Handbook. OBIS project. http://www.obisproject.com (2014-05-29).

Chemla, D., Meunier, F., Wolfler Calvo, R., 2013. Bike sharing systems: Solving the static rebalancing problem. Discrete Optimization 10, 120–146.

Coelho, L.C., Cordeau, J.F., Laporte, G., 2014. Thirty Years of Inventory Routing. Transportation Science 48, 1–19.

Di Gaspero, L., Rendl, A., Urli, T., 2013. A Hybrid ACO+CP for Balancing Bicycle Sharing Systems, in: Hybrid Metaheuristics. Lecture Notes in Computer Science, Springer, pp. 7919:198–212.

Dror, M., Ball, M., Golden, B., 1985. A computational comparison of algorithms for the inventory routing problem. Annals of Operations Research 4.

Gewista Werbegesellschaft m.b.H., 2014. Citybike Wien. http://www.citybikewien.at (2014-06-12).

Kall, P., Wallace, S., 1994. Stochastic Programming. John Wiley & Sons.

Kloimüllner, C., Papazek, P., Hu, B., Raidl, G.R., 2014. Balancing Bicycle Sharing Systems: An Approach for the Dynamic Case, in: Evolutionary Computation in Combinatorial Optimization. Lecture Notes in Computer Science, Springer, pp. 8600:73–84.

Laporte, G., 1992. The vehicle routing problem: An overview of exact and approximate algorithms. European Journal of Operational Research 59, 345–358.

McCarthy, L.M., Knox, P.L., 2005. Urbanization: An introduction to urban geography. Pearson Prentice Hall.

Rainer-Harbach, M., Papazek, P., Hu, B., Raidl, G.R., 2013. Balancing bicycle sharing systems: a variable neighborhood search approach, in: Evolutionary Computation in Combinatorial Optimization, Lecture Notes in Computer Science, Springer. pp. 7832:121–132.

Raviv, T., Tzur, M., Forma, I.A., 2013. Static repositioning in a bike-sharing system: models and solution approaches. EURO Journal on Transportation and Logistics 2, 187–229.

Schuijbroek, J., Hampshire, R., van Hoeve, W.J., 2013. Inventory Rebalancing and Vehicle Routing in Bike Sharing Systems. Tepper School of Business, Paper 1491. http://repository.cmu.edu/tepper/1491 (2014-05-20).

Ulmer, M.W., Brinkmann, J., Mattfeld, D.C., 2015. Anticipatory Planning for Courier, Express and Parcel Services, in: Logistics Management. Lecture Notes in Logistics, Springer.

Vogel, P., Greiser, T., Mattfeld, D.C., 2011. Understanding bike-sharing systems using data mining: exploring activity patterns. Procedia-Social and Behavioral Sciences 20, 514–523.

Vogel, P., Neumann Saavedra, B.A., Mattfeld, D.C., 2014. A hybrid metaheuristic to solve the resource allocation problem in bike sharing systems, in: Hybrid Metaheuristics. Lecture Notes in Computer Science, Springer, pp. 8457:16–29.