Predictive user-based relocation through incentives in one-way car-sharing systems

Short paper submitted to hEART2020

Patrick Stokkink, Nikolas Geroliminis*
Urban Transport Systems Laboratory (LUTS)
École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

February 15, 2020

* Corresponding author
Email addresses: patrick.stokkink@epfl.ch (Patrick Stokkink), nikolas.geroliminis@epfl.ch (Nikolas Geroliminis)

1 Introduction

Due to their benefits in terms of mobility and sustainability, Car-Sharing Systems (CSSs) have become an interesting alternative to private vehicles. Benefits for the individual users include reduced transportation costs and mobility enhancement, while society as a whole benefits from reduced congestion and emissions (see for example Martin and Shaheen (2011) and Baptista et al. (2014)). Over the last years, the number of car-sharing users has showed a rapid increase and a recent study by Frost & Sullivan (2016) has shown that the number of users of CSSs is likely to continue increasing in the future.

We consider a station-based one-way VSS. Station-based VSSs require vehicles to be picked up and dropped off at stations that are located inside the perimeter. One-way systems, unlike two-way systems, allow the customer to drop their vehicle off at any station of their choice. Due to the increased level of flexibility that arises in this system, it is commonly viewed as a more attractive alternative for customers compared to two-way systems.

As described by Boyacı et al. (2015), the attractiveness of a VSS is not only determined by its flexibility, but also by its level of service. The service level is mainly determined by the availability of vehicles at the desired pick-up station and the availability of parking spaces at the desired drop-off station. In one-way vehicle sharing systems, the availability of vehicles is often problematic. Uncertain and asymmetric demand (Jorge and Correia (2013)) are the main causes of the existence of balancing problems. At a station where the demand for vehicles is high, the number of available vehicles declines rapidly. On the other hand, at a station where the supply for vehicles is high, the number of available parking places declines. Due to the limited availability of both vehicles and parking spaces, the level of service of a VSS decreases.

To solve this balancing problem, vehicles should be relocated. The most common type of vehicle relocation is staff-based relocation, where staff members pick up the vehicles at over-saturated locations and deliver them at under-saturated locations. In bike sharing systems, a truck can be used to relocate multiple vehicle at the same time by one staff member (Caggiani and Ottomanelli (2013)). However, in car-sharing systems this procedure is less efficient as only a single car can be moved at the same time.

As an alternative to staff-based relocation, user-based relocation can be used. User-based relocation refers to the case where users are encouraged to relocate the vehicles themselves. A common type of user-based relocation is through customer incentivization, where customers are stimulated to pick up or deliver their vehicle at a different station by offering them a discount. By doing this, the customer helps to reduce the balancing problem. Correia et al. (2014) investigate that if customers are more flexible in their choice for pickup and delivery locations, a significant increase in profit can be obtained by offering incentives to customers. Angelopoulos et al. (2016) provide discounts to customers if they contribute to the balancing process. Their decisions are based on priorities that are assigned based on the capacity and occupancy of the stations.

Most of the literature considers policies where incentivization decisions are made based on threshold values (Clemente et al. (2017)) or problematic scenarios at stations such as being completely full or completely empty (Singla et al. (2015)). These approaches can be classified as non-predictive, in the sense that they do not incorporate expected future demand.

This aim of this paper is to introduce a predictive incentivization policy that determines the optimal incentive based on both the current state of the CSS (i.e. distribution of vehicles throughout the network) and predicted future states, caused by expected future demand. Our strategy uses unobserved customer preferences, which can be learned from previously observed customer behaviour using machine learning algorithms.

By using incentives, we aim to stimulate customers to relocate vehicles from over-saturated to under-saturated stations. In doing so, we can anticipate expected future demand and therefore avoid expected future demand losses. We evaluate our approach using a simulation model with synthetic data from a real experiment, which allows us to compare our methods to existing relocation policies. Our findings show that incentives can increase the service level of the car-sharing system and decrease the number of staff members needed to attain this level. Furthermore, our results indicate that user-based relocations are a more sustainable and profitable way of relocating vehicles compared to staff-based relocations.

The remainder of this paper is organized as follows. Section 2 describes the methodology. Section 3 provides the simulation results and Section 4 summarizes the main conclusions and directions of future research.

2 Methodology

In our framework, we consider and model the operator and the customer separately. It is important to regard the interaction between the operator and the customer. The decision of the customer is influenced by the incentive that is offered by the operator. For the customer, the decision to accept an offer is usually based on the size of the discount, compared to the increase in access time that is experienced. On the other hand, the decision of the operator is influenced by the customer's willingness to cooperate, as well as the arrival rates of future customers.

The characteristics of the designed system are similar to those considered by Repoux et al. (2019) and occur in many real cities such as Toyota City and Grenoble. We consider a one-way car-sharing system where once customers arrive to the system, they select their preferred origin and destination station. The customer is allowed to reserve the vehicle a short time in advance, and a parking space is reserved at the destination until the vehicle is returned. We consider that the customer is possibly offered an alternative and less convenient trip (as it is not the customer's first preference, the access time increases) for a lower price. Due to shortage of either vehicles or parking spaces, it is possible that a customer's first choice trip is unavailable. In that case, we assume that the customer can either choose to accept the incentive or decline the incentive and choose a different mode of transportation.

2.1 Operator Decision Model

We define I as the set of possible incentives. For every incentive i we define $\Delta_{time}(i)$, the increase in access time the customer experiences when accepting the incentive, and $\Delta_{costs}(i)$, the decrease in price experienced by accepting the incentive. The estimated probability that a customer accepts an incentive, $\hat{P}_{acc}$, is based both on $\Delta_{time}(i)$ and $\Delta_{costs}(i)$. The shape of the estimated probability function $\hat{P}_{acc}(\Delta_{time}, \Delta_{costs})$ is described in more detail in Section 2.2.

The aim of offering incentives is to relocate the vehicles in order to omit future losses in demand due to vehicle imbalances. We refer to this as the expected omitted demand loss, $ODL(i)$, experienced when the customer accepts incentive $i \in I$. Thereby, we define $\bar{c}$ as the average profit obtained from one unit of demand. We refer to the original pickup and delivery stations as o and d respectively. The pickup and delivery stations that are chosen as a consequence of the acceptance of the incentive are referred to as o^* and d^*. As depicted in Figure 1, the use of one incentive implicitly replaces at most two staff-based relocations.

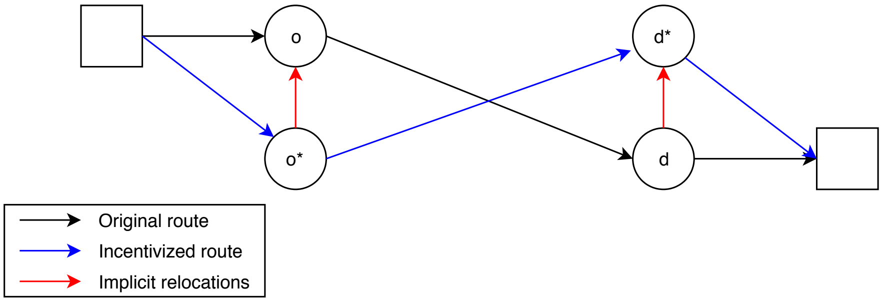

Figure 1: Implicit relocations experienced due to incentive

The operator can determine for every possible incentive $i \in I$ what the optimal discount $\Delta_{costs}(i)$ is. This decision is based on a trade-off between costs and the probability with which the incentive is accepted. Note that for a given incentive i, the corresponding value of $\Delta_{time}(i)$ is fixed. The objective can be formulated as follows:

$$f_i(\Delta_{costs}(i)) = \max_{\Delta_{costs}(i)} \hat{P}_{acc}(\Delta_{time}(i), \Delta_{costs}(i)) \cdot (ODL(i)\bar{c} - \Delta_{costs}(i)) \tag{1}$$

The objective is to maximize the expected additional profit by offering the incentive. This additional profit consists of the profit obtained by the omitted demand loss, minus the discount that was offered. The optimal value of each incentive can be found by solving the optimization problem in (1). Thereafter, the best incentive can be chosen by optimizing over the set of possible incentives. We refer to the optimal incentive as i^* and to the corresponding optimal value of the incentive as $\Delta^*_{costs}(i^*)$ (which is the argument of the subproblem). The total objective can therefore be formulated as follows:

$$i^* = \arg\max_{i \in I} f_i(\Delta^*_{costs}(i)) \tag{2}$$

To determine the value of $ODL(i)$ we construct a Markovian model expanding the model proposed by Repoux et al. (2019). We consider a separate Markov chain for every station, which allows us to define the expected loss of future customer demand given the current state of the system. Loss of customer demand is encountered if either the desired pickup station has no available vehicles or the desired drop-off station has no available parking spaces. Stations can be selected as either origin or destination locations of relocations. The omitted demand loss for these stations is defined as O_s and D_s respectively. By selecting a station to be an origin location, the number of available vehicles decreases by one, which is reserved for the customer that will relocate it later. At a destination location, a parking space is reserved for this same customer. The omitted demand loss at each station can then be calculated as the difference between the expected demand loss in the old and the new situation.

We define $ODL(i)$ by considering the implicit relocations described in Figure 1. Stations o^* and d are selected as origin stations of the relocation, and stations o and d^* are destinations of the relocation. Therefore, the total omitted demand loss is defined as follows:

$$ODL(i) = (O_{o^*} + D_o) + (O_d + D_{d^*}) \tag{3}$$

2.2 Customer Decision Model

To model the customer decisions, we define the probability function P_{acc}. We assume that the probability with which a customer accepts an incentive depends on the value of the incentive and the additional access time experienced because of this incentive. This type of customer decisions is commonly modelled using a logistic model. A similar model has been used by Di Febbraro et al. (2018). The acceptance probability is defined as follows:

$$P_{acc} = \frac{1}{1 + e^{-(\beta X)}} \tag{4}$$

with $X = [1,\ \Delta_{time},\ \Delta_{costs}]$. Two important notes have to be made considering this probability function. First, customers are heterogeneous in the sense that they value time differently. This suggests that the parameter β is customer-specific. Second, the actual shape of the probability function P_{acc} is unknown to the operator. The operator can however use previously observed data to create an estimation of the parameters $\hat{\beta}$ through learning over time and thereby estimate the probability function $\hat{P}_{acc}$.

A customer's value of time can be determined as the relative importance of the coefficients β_{time} and β_{costs}, which can be estimated by taking the ratio of the two. The higher the ratio of these coefficients, the higher a customer's value of time. A customer's value of time can be interpreted as the additional discount a customer wishes to receive for one minute of extra access time.

In addition to this probability model, some hard constraints may apply to the choice to accept an incentive. We apply one hard constraint, which says that an incentive is only accepted if the additional access time is at most 10 minutes to the pickup station or from the delivery station. This constraint can be easily incorporated in the optimization problem in (2).

The acceptance probability function $\hat{P}_{acc}$ can be estimated using all previously observed data. That is, by storing the information of the observed incentive (X) and the customer's response to the offered incentive (accept or reject), we can estimate the coefficients β using Maximum Likelihood (ML).

3 Simulation Results

Table 1 presents the simulation results under the assumption that the operator has perfect information on the value of time of customers. The results are generated using 100 simulation runs of 10 consecutive days. The profit is based on various costs as defined by Boyacı et al. (2015). The profit is calculated as the user revenue based on a price of €0.20 per minute minus the costs of relocators (€18 per hour), fixed vehicle costs (€20 per day) and costs per kilometer travelled by both users and relocators (€0.01 per KM). We use instances based on the Grenoble carsharing system, as introduced by Repoux et al. (2019). We consider the case where daily demand is assumed to be equal to 100 customers and 40 vehicles are used throughout the network, which consists of 28 stations. We compare our incentivization policy with the Markovian staff-based relocation policy as described by Repoux et al. (2019).

Table 1: Simulation results

Max time	Staff	Incentives	% served	# relocations	# incentives	KM travelled	Profit
8	0	No	59.7	0.0	0.0	5.82	100.00
		Yes	64.7	0.0	9.3	5.75	165.20
	1	No	84.4	31.9	0.0	6.34	316.81
		Yes	86.1	30.8	13.9	6.32	297.65
	2	No	89.0	58.5	0.0	6.94	73.48
		Yes	90.4	57.7	15.5	6.91	46.57
	3	No	90.4	78.8	0.0	7.40	-236.93
10	0	No	59.7	0.0	0.0	5.82	100.00
		Yes	66.6	0.0	11.1	5.72	184.22
	1	No	84.4	31.9	0.0	6.34	316.81
		Yes	87.0	30.3	16.2	6.31	290.31
	2	No	89.0	58.5	0.0	6.94	73.48
		Yes	**90.8**	57.4	11.7	**6.90**	**31.76**
	3	No	**90.4**	78.8	0.0	**7.40**	**-236.93**
15	0	No	59.7	0.0	0.0	5.82	100.00
		Yes	74.3	0.0	19.3	5.69	239.13
	1	No	84.4	31.9	0.0	6.34	316.81
		Yes	**89.4**	29.0	24.6	**6.34**	**230.12**
	1	No	**89.0**	58.5	0.0	**6.94**	**73.48**
		Yes	92.0	57.1	24.2	6.88	-41.67
	3	No	90.4	78.8	0.0	7.40	-236.93

The first column denotes the maximum extra access time. The following two columns describe the relocation policy that is used (i.e. number of staff members and whether incentives are used). The fourth column denotes the percentage of served customers. The fifth and sixth column are a daily average of the number of relocations performed and the number of incentives accepted respectively. The KM travelled is measured as an average per served demand unit and includes both the user and staff KM travelled. The profit is given as a percentage relative to the scenario without staff and user relocations.

The results indicate that the use of incentives is beneficial in terms of the percentage of served demand. From the average number of relocations per day we observe that the use of incentives does not replace the staff-based relocations, but rather generates additional improvements to the system. Although incentives lead to the generation of less additional demand compared to staff-based relocations, incentives are also substantially less costly.

An interesting comparison can be made between the case where two staff members and incentives are used and the case where three staff members and no incentives are used, under the assumption that the maximum extra access time is equal to 10 minutes. As these cases obtain the same service rate, we can make a fair comparison in terms of the other components. First of all, we observe that for every demand unit, the average distance travelled is approximately 500 meters lower. As approximately 90 demand unites are served per day, this constitutes to a decrease of almost 50 KM per day. User-based relocations are therefore more sustainable than staff-based relocations. In addition to this, a comparison of the profit of those scenarios shows that replacing the third staff member by incentives, has a significant positive effect of the earned profit.

We observe that as the maximum access time increases, more beneficial incentives can be found.

However, we note that it is unlikely that customers are willing to walk for 15 minutes to pickup their vehicle. A potential solution for this is to allow customers to use public transport to reach their vehicle. As a direction of further research, we aim to integrate car-sharing systems with public transport to reduce the access time and thereby increase the effectiveness of incentives.

3.1 Adapted staff-based relocation policy

In the staff-based relocation policy defined by Repoux et al. (2019), a staff-member executes any relocation job with positive ODL. We can adapt the staff-based relocation policy to the use of incentives by introducing a threshold value τ, such that a relocation is only executed if the ODL is larger than τ. By changing this value we can reduce the number of relocations performed by staff members and thereby reduce their overall activity. This can be convenient as in practice, staff members often perform maintenance jobs if they are not relocating vehicles. For this experiment, we assume the costs of staff-members are proportional to their activity. We consider various threshold values varying between 0.00 and 0.25. All other simulation settings are similar to those of the previous experiment and a maximum access time of 10 minutes is used. The results of this experiment are provided in Table 2.

Table 2: Simulation results for different staff thresholds

Incentives	τ	% served	# relocations	# incentives	KM travelled	Profit	% active
No	0.00	84.4	31.9	0.0	6.34	100.00	98.5
Yes	0.00	87.0	30.3	16.2	6.31	92.18	98.2
	0.10	86.8	27.5	16.7	6.28	93.80	94.4
	0.15	86.2	23.9	17.1	6.23	97.24	86.3
	0.20	85.0	19.6	17.3	6.15	100.59	74.7
	0.25	83.1	15.5	16.8	6.06	101.91	60.9

We consider the case where one staff member is used. The first two columns describe the scenario. The third column denotes the percentage of served customers. The fourth and fifth column are a daily average of the number of relocations performed and the number of incentives accepted respectively. The KM travelled is measured as an average per served demand unit and includes both the user and staff KM travelled. The profit is given as a percentage relative to the scenario without user relocations. The last column denotes the percentage of time the relocator is active

The results indicate that by increasing the threshold value, we reduce the number of relocations and the percentage of time the staff member is relocating vehicles. This suggests that we can reduce the activity of staff-members up to 25 to 40%, while achieving a similar profit and service rate. We also observe that the number of accepted incentives increases with the threshold, as part of the relocations that have an ODL value beneath the threshold are now executed by customers. This effect is no longer visible for $\tau > 0.20$, which can be clarified by looking at the distribution of the ODL as depicted in Figure 2. This figure indicates that the ODL obtained from incentives is rarely higher than 0.20. Intuitively, this result follows from the observation that user-based relocations cover a relatively shorter distance than staff-based relocations, as users are not willing to walk too far. As the imbalances among stations within a small radius are relatively small, the ODL from a relocation between those stations is also smaller.

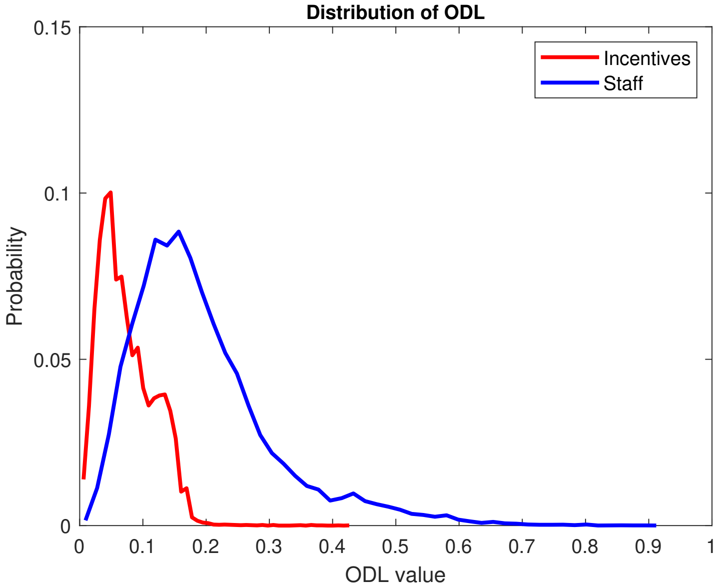

Figure 2: Distribution of ODL by incentive and staff relocation

3.2 Learning from customer behaviour

In general, the value of time of customers is unknown to the operator. Therefore, we investigate the effect of the use of our learning method to estimate the customers' average value of time, opposed to the case where the exact value of time of customers is assumed to be known. The results of this experiment are displayed in Table 3 and are simulated without staff-based relocations.

Table 3: Simulation results when learning method is used

Incentives	Learning	% served	% accepted	Discount value	Profit
No	No	59.7			100.00
Yes	Yes	65.6	74.7	38.64	169.44
Yes	Perfect Info	66.6	98.9	36.57	184.22

The first two columns describe the settings of the simulation. The third and fourth column denote the percentage of served customers and accepted incentives respectively. The average discount value is in cents per additional minute travelled. The profit is given as a percentage relative to the scenario without staff and user relocations.

We observe that without perfect information on the value of time of customers, a smaller portion of the offered incentives is accepted. As a consequence, the service rate and profit are slightly lower. However, we observe that by using incentives in combination with our learning method a significant improvement over the scenario without relocations is made in terms of both service rate and profit.

We note that the quality of our learning algorithm is dependent on the distribution of the value of time of customers. If large differences exist between customer segments, the performance of our algorithm decreases. Therefore, as a direction of further research, we aim to develop a clustering algorithm that groups customers with similar behaviour, such that cluster-specific estimations can be obtained.

4 Conclusion

In this paper, we proposed a new predictive user-based vehicle relocation policy for one-way car-sharing systems. The policy uses information on the current state of the system, as well as expected future demand and estimated customer preferences to determine the optimal user-based relocation and corresponding discount value (incentive). Our results indicate that the proposed user-based relocations are a cheap and sustainable way to relocate vehicles. Especially when our policy is combined with a staff-based relocation policy, significant improvements in service level, profit and the number of kilometers travelled can be made.

The main directions of further research include the development of a clustering algorithm to group customers with similar behaviour and thereby improve the learning procedure. In addition to this, we aim to integrate car-sharing systems with public transport to reduce the access time to available vehicles and thereby increase the effectiveness of incentives.

References

Angelopoulos, A., Gavalas, D., Konstantopoulos, C., Kypriadis, D., and Pantziou, G. (2016). Incentivization schemes for vehicle allocation in one-way vehicle sharing systems. In *2016 IEEE International Smart Cities Conference (ISC2)*, pages 1–7. IEEE.

Baptista, P., Melo, S., and Rolim, C. (2014). Energy, environmental and mobility impacts of car-sharing systems. empirical results from Lisbon, Portugal. *Procedia-Social and Behavioral Sciences*, 111:28–37.

Boyacı, B., Zografos, K. G., and Geroliminis, N. (2015). An optimization framework for the development of efficient one-way car-sharing systems. *European Journal of Operational Research*, 240(3):718–733.

Caggiani, L. and Ottomanelli, M. (2013). A dynamic simulation based model for optimal fleet repositioning in bike-sharing systems. *Procedia-Social and Behavioral Sciences*, 87:203–210.

Clemente, M., Fanti, M. P., Iacobellis, G., Nolich, M., and Ukovich, W. (2017). A decision support system for user-based vehicle relocation in car sharing systems. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 48(8):1283–1296.

Correia, G. H. D. A., Jorge, D. R., and Antunes, D. M. (2014). The added value of accounting for users' flexibility and information on the potential of a station-based one-way car-sharing system: An application in lisbon, portugal. *Journal of Intelligent Transportation Systems*, 18(3):299–308.

Di Febbraro, A., Sacco, N., and Saeednia, M. (2018). One-way car-sharing profit maximization by means of user-based vehicle relocation. *IEEE Transactions on Intelligent Transportation Systems*, 20(2):628–641.

Frost & Sullivan (2016). Future of carsharing market to 2025. *Technology Advancements, Market Consolidation and Government Initiatives to Influence Market Growth Over the Next Decade.*

Jorge, D. and Correia, G. (2013). Carsharing systems demand estimation and defined operations: a literature review. *European Journal of Transport and Infrastructure Research*, 13(3).

Martin, E. and Shaheen, S. (2011). The impact of carsharing on public transit and non-motorized travel: an exploration of north american carsharing survey data. *Energies*, 4(11):2094–2114.

Repoux, M., Kaspi, M., Boyacı, B., and Geroliminis, N. (2019). Dynamic prediction-based relocation policies in one-way station-based carsharing systems with complete journey reservations. *Transportation Research Part B: Methodological*, 130:82–104.

Singla, A., Santoni, M., Bartók, G., Mukerji, P., Meenen, M., and Krause, A. (2015). Incentivizing users for balancing bike sharing systems. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.