

Incentivizing Users for Balancing Bike Sharing Systems

Adish Singla

ETH Zurich
adish.singla@inf.ethz.ch

Marco Santoni

ElectricFeel Mobility Systems
marco@electricfeel.com

Gábor Bartók

ETH Zurich
bartok@inf.ethz.ch

Pratik Mukerji

ElectricFeel Mobility Systems
mukerji@electricfeel.com

Moritz Meenen

ElectricFeel Mobility Systems
meenen@electricfeel.com

Andreas Krause

ETH Zurich
krausea@ethz.ch

Abstract

Bike sharing systems have been recently adopted by a growing number of cities as a new means of transportation offering citizens a flexible, fast and green alternative for mobility. Users can pick up or drop off the bicycles at a station of their choice without prior notice or time planning. This increased flexibility comes with the challenge of unpredictable and fluctuating demand as well as irregular flow patterns of the bikes. As a result, these systems can incur imbalance problems such as the unavailability of bikes or parking docks at stations. In this light, operators deploy fleets of vehicles which redistribute the bikes in order to guarantee a desirable service level. Can we engage the users themselves to solve the imbalance problem in bike sharing systems?

In this paper, we address this question and present a crowdsourcing mechanism that incentivizes the users in the bike repositioning process by providing them with alternate choices to pick or return bikes in exchange for monetary incentives. We design the complete architecture of the incentives system which employs optimal pricing policies using the approach of regret minimization in online learning. We investigate the incentive compatibility of our mechanism and extensively evaluate it through simulations based on data collected via a survey study. Finally, we deployed the proposed system through a smartphone app among users of a large scale bike sharing system operated by a public transport company, and we provide results from this experimental deployment. To our knowledge, this is the first dynamic incentives system for bikes re-distribution ever deployed in a real-world bike sharing system.

Introduction

A Bike Sharing System (henceforth BSS) is a new concept of public transportation offering citizens a flexible and green alternative as well as complementing the slow and crowded transportation in urban areas. This mobility trend had an exponential growth over the last years, and, as of June 2014, over 700 cities actively operate automated bike sharing systems deploying an estimate of 800,000 bikes worldwide (Meddin and DeMaio 2014). BSSs increase the user's commuting flexibility by allowing her to pick up or drop off

a bike at any station and let her decide the duration of the trip without any prior planning or reservation. Indeed, this flexibility is the key factor for the success of BSSs.

Balancing problem. This flexibility also poses a number of new challenges for the BSS operators. The demand is often unpredictable, asymmetric and fluctuating throughout the day. Other factors such as altitude differences, weather conditions or events in the city can cause irregular or asymmetric rental demands and flow of bikes. As the number of available bikes and parking spots at any bike station in the system is limited, satisfying the forthcoming demand with such limited resources is a major challenge and recurrent problem for BSS operators. A BSS user would experience an unsatisfactory quality of service when attempting to rent or drop off a bike at an empty or full station respectively. Meeting the demand for bikes and free docks requires an extensive logistic work and poses major operational costs to the operators. Thus, fleets of vehicles (or trucks) are deployed to re-distribute bikes among stations to balance the system out (Büttner, Mlasowsky, and Birkholz 2011).

Incentivizing users. Apart from the large operational costs, deploying trucks for the re-distribution goes against the green concept of BSS. Can we engage the users of such systems and provide them incentives to contribute rebalancing the system? Termed as *crowdphysics* by Sadilek, Krumm, and Horvitz (2013), this class of crowdsourcing systems requires users to sequence or synchronize physical actions in time and space, and can lead to building self-sustainable systems. The goal of building such a system presents several challenges including: *i*) the system's large scale in terms of user base as well as number of stations, *ii*) the fluctuating demand over time that needs to be taken into account while determining problematic stations, *iii*) the unknown personal costs of the users and their strategic behavior potentially aiming at maximizing their profit, *iv*) the budget constraints of how much can be spent on incentives and *v*) a user-friendly interface to make such a system appealing (e.g., through a smartphone app). Building such a self-sustainable and crowd-powered bike sharing system is the key idea we explore in this paper. While we focus on BSSs, similar problems arise in other domains (e.g., car sharing or rerouting users on over-booked flights) and our methodology is applicable to these systems as well.

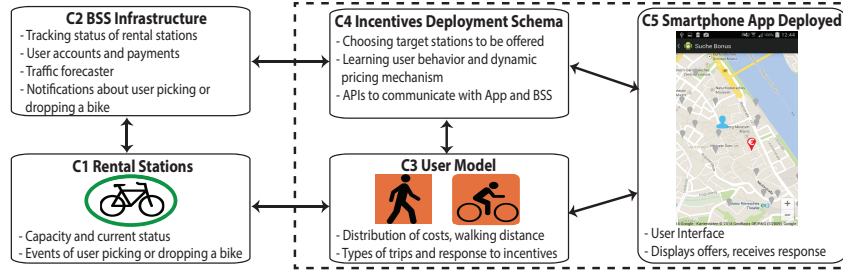


Figure 1: Overview of the system, our contributions are shown in the dotted box

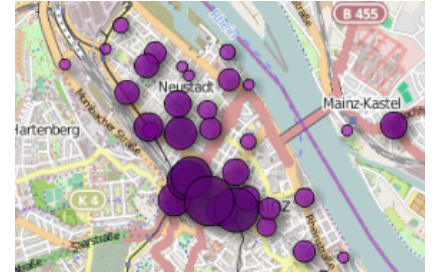


Figure 2: Spatial dist. of accepted offers

Overview of our system

In this paper, we design the complete architecture of an operating incentives system for engaging users in the bike-repositioning effort (see Figure 1). In collaboration with a company operating large scale bike sharing systems, we deployed our system on a real-world BSS in a city of Europe¹. In Figure 1, the components *C1 Rental Stations* and *C2 BSS Infrastructure* correspond to the rental stations, the hardware and software infrastructure of the public bike sharing system. These components involve tracking all the user activities such as picking or dropping the bikes, managing user accounts and the payment system.

The main components that we designed and built include *C3 User Model*, *C4 Incentives Deployment Schema* (IDS) and *C5 Smartphone App*. The IDS is the central component in the overall design that handles the user’s requests through *C5 Smartphone App*. The request includes input such as the user’s target station and type (i.e., bike pick up or drop off). The IDS communicates with the BSS Infrastructure to evaluate the current and predicted status of the stations, and then decides whether to offer incentives to the user by requesting a change in her pickup or drop-off location. In order to maximize the efficiency under given budget constraints, we design dynamic pricing mechanisms using the approach of regret minimization in online learning that can learn over time about the optimal pricing policies. We consider the users as strategic agents who may untruthfully report information about their personal cost and location to maximize their profit. We model the full rental process of the user and her reaction to incentives in the component *C3 User Model*. In summary, our main contributions are as follows:

- We design the complete architecture of an incentive system with the goal of enabling self-sustainable, crowd-powered and greener bike sharing systems.
- We build an incentive deployment schema that interacts with the users and BSS infrastructure, decides about offering incentives, and learns from past interactions.
- We validate our assumptions by surveying BSS users, and demonstrate the effectiveness of our approach via a detailed BSS simulator. We report on our experience deploying our system via a smartphone app integrated with a real-world BSS.

Related Work

Truck-based repositioning in BSS. The problem of asymmetric flow, imbalance and repositing of the vehicles has

been studied in generic vehicle sharing systems that could involve cars or bikes (Cepolina and Farina 2012; Kek et al. 2009; Clemente et al. 2013; Nair and Miller-Hooks 2011; Correia and Antunes 2012). Regarding BSS’s balancing problem, a substantial literature (Rainer-Harbach et al. 2013; Pfrommer et al. 2014; Raviv and Kolka 2013; Chemla, Meunier, and Wolfier Calvo 2013; Schuijbroek, Hampshire, and van Hoeve 2013; George and Xia 2011) introduced bike repositioning policies based on trucks. The user dissatisfaction function introduced by Raviv and Kolka (2013) measures the performance of a BSS station and provides the quality of a repositioning. Nair et al. (2013) provides a detailed quantitative analysis of the repositioning strategies for Vélib’, the BSS in Paris.

Crowd-based repositioning in BSS. Recent literature has investigated crowd-based BSS balancing policies (Fricker and Gast 2012; Wasserhole, Jost, and Brauner 2013; Pfrommer et al. 2014; Chemla, Meunier, and Pradeau 2013; GoDCgo 2011) to study the potential of influencing BSS users by offering incentives that could improve balancing of the system. In Paris, Vélib’ runs a static incentive schema that offers users 15 extra free minutes each time they return a bike to an elevated station (Vélib’ 2014). However, in contrast, our proposed mechanism is dynamic and can efficiently take into account the current state of the system to decide where and what incentives to offer. Pfrommer et al. (2014) also introduced a dynamic pricing mechanism using model-based predictive control principles. Our pricing mechanism, in contrast, is based on an efficient and provably near-optimal online learning framework. Furthermore, we build an operating end-to-end system integrated with a real-world BSS for deployment, and we investigate the incentive compatibility of our approach.

The Model

We now formalize the problem addressed in this paper.

BSS. There are m stations in a city denoted by the set $S = \{s_1, s_2, \dots, s_m\}$. We denote the number of bikes available at time t at station s as $v(s, t)$. Along the lines of Pfrommer et al. (2014), we split every day into 12 slices $h \in \mathcal{H}$ of two hours each, where a time of the day t is mapped to the corresponding time-slice using $h(t)$. The demand of rentals from station s_i to station s_j within time-slice h is modeled as a discrete random variable z which represents a count of the number of trips with probability density function $\tau(z|s_i, s_j, h)$. We denote $\hat{z}(h)$ to represent the estimate of the total number of trips within time-slice h . The BSS infrastructure has a proprietary demand forecaster that can

¹Bike sharing operator MVGmeinRad in Mainz, Germany.

predict the future demand of incoming and outgoing traffic at the stations and is made available to us through their APIs. See (Côme and Oukhellou 2012; Kaltenbrunner et al. 2010; Borgnat et al. 2011; Froehlich, Neumann, and Oliver 2009; Han, Côme, and Oukhellou 2014) for possible approaches.

Quality of service. The goal of a bike repositioning policy is to prevent customers' dissatisfaction about not finding an available bike or a parking slot. Waserhole, Jost, and Brauner; Chemla, Meunier, and Pradeau; Fricker and Gast (2013; 2013; 2012) have proposed metrics aiming at maximizing the number of trips or minimizing the number of problematic stations. However, these metrics fail to capture the dissatisfaction experienced by BSS users. We adopt the measure of service level proposed in Pfrommer et al. (2014):

$$\text{Service level} = \frac{\text{Potential customers} - \text{No-service events}}{\text{Potential customers}}$$

where the number of no-service events corresponds to the customers who could not rent a bike at an empty station or failed to return their bike at a full station.

Truck-based re-positioning policies. The BSS operator allocates a fixed daily budget B that can be used by the truck-based policy. The BSS company where we deployed our incentives system uses its proprietary re-positioning policy. See (Chemla, Meunier, and Pradeau 2013; Raviv and Kolka 2013) for possible approaches.

User model. Consider a user (or customer) u at location l_x^u interested in an action $x \in \{\text{pick}, \text{return}\}$, i.e., pick up or drop off a bike. We assume that, by default, the user would consider the nearest station to her location to perform the action, given by $s_x^u \leftarrow \arg \min_{s \in S} d(l_x^u, s)$, where d measures the distance between two locations. The incentives system encourages users to shift their pickup or drop-off stations to contribute balancing the BSS. For simplicity, we assume that the user u is willing to walk to another station up to a maximum distance γ^u , and, within this radius, her cost for the additional effort is constant given by c^u . She is not willing to walk for distance more than γ^u despite of any offer by our incentive system. We validate these assumptions through a survey study with real BSS customers. In our model, the costs of the users are *i.i.d.* sampled from an underlying *unknown* distribution f . We let $F : [c_{\min}, c_{\max}] \rightarrow [0, 1]$ denote the cumulative distribution function (CDF) of costs or "cost curve", unknown to the system.

Suppose a user u at location l_x^u , going to station s_x^u , is offered an alternate station s_x^* for a payment of p^* . We model her reaction to this incentive as an indicator function:

$$\mathbb{1}_{\text{accept}}^u(l_x^u, s_x^*, p^*) = \begin{cases} 0 & \text{if } \gamma^u < d(l_x^u, s_x^*) \text{ or } p^* < c^u \\ 1 & \text{otherwise} \end{cases}$$

which takes value 0 or 1 when the user rejects or accepts the offer, respectively. We remark that the information about user cost c^u and her true location l_x^u is private to the user and may vary among users. We consider the users as strategic agents who may untruthfully report their private information about the location l_x^u or private cost c^u to maximize their profit. For simplicity of the mechanism design, we use $\hat{\gamma}$ to represent the maximum distance for the whole population, and we obtain its estimate through a user survey.

Procedure 1: Incentives Deployment Schema (IDS)

```

1 Input: user  $u$ ; location  $l_x^u$ ;  $x \in \{\text{pick}, \text{return}\}$ ;
2 Output: Offer: station  $s_x^*$ ; price  $p^*$ ;
3 Parameters: radius  $\hat{\gamma}$ ;
begin
4   if HASINCENTIVESPENDING( $u$ ) then
5     return null;
6   candidate stations  $C \leftarrow S$ ;
7   default target  $s_x^u \leftarrow \arg \min_{s \in S} d(l_x^u, s)$ ;
8   foreach  $s \in C$  do
9      $s.\text{status} \leftarrow \text{PROBLEMATIC}(s)$ ;
10    if  $d(l_x^u, s) > \hat{\gamma}$  then
11       $C = C \setminus \{s\}$ ;
12    if  $x = \text{return}$  &  $s.\text{status} \neq \text{EMPTY}$  then
13       $C = C \setminus \{s\}$ ;
14    if  $x = \text{pick}$  &  $s.\text{status} \neq \text{FULL}$  then
15       $C = C \setminus \{s\}$ ;
16  if  $s_x^u \in C$  OR  $C = \emptyset$  then
17    return null;
18   $s_x^* \leftarrow \arg \min_{s \in C} d(l_x^u, s)$ ;
19   $p^* \leftarrow \text{PRICINGMECHANISM}()$ ;
20  Output:  $s_x^*, p^*$ ;

```

Incentivizing users for repositioning. We study an incentive system as an alternative to truck-based re-distribution in BSSs. We propose a dynamic incentives schema that computes incentives each time a user makes a new request. The BSS operator regularly allocates a budget $B(h)$ for time batch h , and the goal of the incentives schema is to exploit the budget efficiently. Therefore, if a user u makes a request at time t for action type x , a dynamic algorithm is responsible to select the alternate station s_x^* and compute the value of the incentive p^* based on the following inputs: *i*) reported location of user l_x^u , *ii*) current status of the system: $v(s, t) \forall s \in S$, *iii*) near-future demand $\tau(z|s_i, s_j, h) \forall s_i, s_j \in S$, *iv*) budget available, and *v*) information about historical interactions with users.

Our goal is to design a system that is *i*) budget-feasible, i.e., operates under strict budget constraints, *ii*) efficient in terms of improvement in quality of service for a given budget, *iii*) incentive-compatible (truthful), i.e., it is in best interest of users to reveal true information, and *iv*) can deal with unknown user costs and learn pricing policies over time.

The Incentives System

We now present the various components of our incentives system and discuss its properties.

Incentives Deployment Schema

Procedure 1 illustrates the *Incentives Deployment Schema* (IDS) that handles the users' requests forwarded by the deployed smartphone app. For each request, the IDS gets as input the user identifier u , the action type $x \in \{\text{pick}, \text{return}\}$ and the user's reported location l_x^u . Based on the user input and the current status of the system, IDS computes the offer that consists of the alternative target station s_x^* and the offered price p^* . Initially, the IDS ensures that the user does

not have any pending incentive, i.e., an incentive offer that she recently accepted but did not accomplish yet. Then it creates a set of candidate stations that are within the maximum walking distance $\hat{\gamma}$ from l_x^u and are “problematic”, i.e. full or empty. The problematic status of the stations is based on input from the APIs of the BSS infrastructure that provides the current status $v(s, t)$ as well as the forecast of future traffic. The IDS also takes into account the pending incentives associated with each station while updating the status that additionally accounts for future traffic.

Once the IDS has filtered the candidate stations, it selects as target station s_x^* the one closest to reported location l_x^u . Then, it sends a request to the *pricing mechanism* to obtain the price p^* to be offered. This offer is displayed to the user on the smartphone app. If the user accepts, the IDS registers this incentive offer as pending until she accomplishes the assigned task or it expires after a prefixed amount of time.

Pricing Mechanism DBP-UCB

The main challenges in designing an optimal pricing mechanism are due to the limited and dynamically changing budget (B) as well as pool size of users to make offers to (N) and the unknown “cost curve” F . Additionally, the mechanism has to compute the offers in an online setting with only limited interaction with the users in terms of the information that can be received from them for learning. First, let us consider an offline setting where the cost curve F is known and $\{p_0 (= c_{min}), \dots, p_i, \dots, p_K (= c_{max})\}$ be the set of available prices that the mechanism can offer. In this case, the optimal truthful mechanism denoted by OPT-FIX, is to offer a fixed price p^{OPT} (Singer 2010) given by:

$$p^{OPT} = \arg \max_p \min \left\{ F(p), \frac{B}{N \cdot p} \right\} \text{ s.t. } p \in \{p_0, \dots, p_K\} \quad (1)$$

Now, in case of unknown F , we can cast the problem as that of online learning of the optimal price. Note that, $F(p)$ simply denotes the probability of getting an acceptance from the user when offered price p . In particular, we are interested in designing mechanisms based on the posted-price model where the mechanism makes a price offer p and the user simply says “yes” or “no” to the offer, instead of soliciting the user’s bid. In a learning framework, the mechanism can make different price offers over time to different users, and use the binary feedback of acceptance to learn an estimate of $F(p)$. There is an inherent explore-exploit trade off here of exploring (or learning) about different, potentially suboptimal, prices and exploiting the estimates of $F(p)$ to offer the price that appears best. Multi-armed bandits (MAB) (Lai and Robbins 1985) provide a natural formalism for tackling this problem, considering the prices as actions (or arms) with unknown stochastic reward given by the function $F(p)$. In the MAB framework, the performance of a mechanism is measured in terms of its “regret” R_n , i.e., the cumulative loss of the mechanism w.r.t. the optimal fixed price p^{OPT} after n iterations. The goal is to design a “no-regret” mechanism that learns to perform competitively, i.e., the average regret w.r.t. n goes to zero $\lim_{n \rightarrow \infty} R_n/n = 0$. However, note that in contrast to standard MAB settings, here actions have dif-

Procedure 2: Pricing Mechanism DBP-UCB

```

1 Input: start time  $t_0$ ; prices  $\{p_0, \dots, p_i, \dots, p_K\}$ ;
2 Output: Price  $p^n$  at iteration  $n$ ;
3 Initialize:
    •  $n = 0$ ;  $h^0 = h(t_0)$ ; ▷ First batch
    •  $B = B(h^n)$ ;  $B^n = B$ ;  $N = 2 \cdot \hat{z}(h^n)$ ;
    •  $N_i^n = 0$ ;  $F_i^n = 0$ ;  $\forall i \in [0, \dots, K]$  ▷ Value estimates
4 foreach request by user at time  $t$  do
5   if  $h^n \neq h(t)$  then
6      $h^n = h(t)$ ; ▷ New time batch
7      $B^n = B^n + B(h^n)$ ;  $B = B^n$ ; ▷ More budget
8      $N = 2 \cdot \hat{z}(h^n)$ ; ▷ Trips estimate from forecaster
9      $\tilde{F}_i^n = F_i^n + \sqrt{\frac{2 \cdot \ln(n)}{N_i^n}}$ ; ▷ Same as in BP-UCB
10     $i^n = \arg \max_i \min \left\{ \tilde{F}_i^n, \frac{B}{N \cdot p_i} \right\} \text{ s.t. } p_i \leq B^n$ ;
11  Output: Offer price  $p^n = p_{i^n}$  at iteration  $n$ ;
12  Feedback: Observe acceptance decision  $y^n$ ;
13  Update Variables:
    •  $B^{n+1} = B^n - p^n \cdot y^n$ ;  $F_{i^n}^{n+1} = F_{i^n}^n + \frac{(y^n - F_{i^n}^n)}{(N_{i^n}^n + 1)}$ ;
    •  $N_{i^n}^{n+1} = N_{i^n}^n + 1$ ;  $h^{n+1} = h^n$ ;  $n = n + 1$ ;
14  Wait for new user request;

```

ferent costs and there is a budget constraint.

For static settings where the budget (B) and the number of users (N) is known and fixed in advance, Singla and Krause (2013) designed a no-regret mechanism BP-UCB that is budget feasible, truthful and has fast convergence rates. This scenario does not correspond to the incentives system we are interested in as, in this context, the pricing mechanism has to operate continuously throughout time, is allocated new budget in time batches and the number of users coming to the system is dynamic. In this light, we design our mechanism DBP-UCB (Dynamic Budgeted Procurement using Upper Confidence Bounds), illustrated in Procedure 2, as a dynamic variant of BP-UCB. The key idea we use is that the learning process of the mechanism can be decoupled from the constraints on B and N . While the parameters B and N change dynamically at the onset of each time batch, it carries over the estimates of $F(p)$ across batches, ensuring DBP-UCB maintains the convergence properties of BP-UCB.

At a high level, the mechanism DBP-UCB operates as follows. At the onset of each new time batch h , the mechanism is provided with an additional budget $B(h)$ by the BSS operator. Furthermore, the number of participants N for a batch is approximated by the expected number of trips $\hat{z}(h)$ taking place in the corresponding batch h estimated by the BSS forecaster. In particular, the expected number of rentals is multiplied by two as the user may interact with the incentives system both to pick and return a bike. The mechanism sequentially interacts with users in discrete steps denoted by n . Let B^n be the budget at iteration n , N_i^n be the number of times p_i price has been offered so far, and F_i^n be the current estimate of the cost curve for prices $\forall i \in [0, \dots, K]$. The mechanism also maintains upper confidence bounds on F_i^n denoted as \tilde{F}_i^n , representing the optimistic estimates of F at

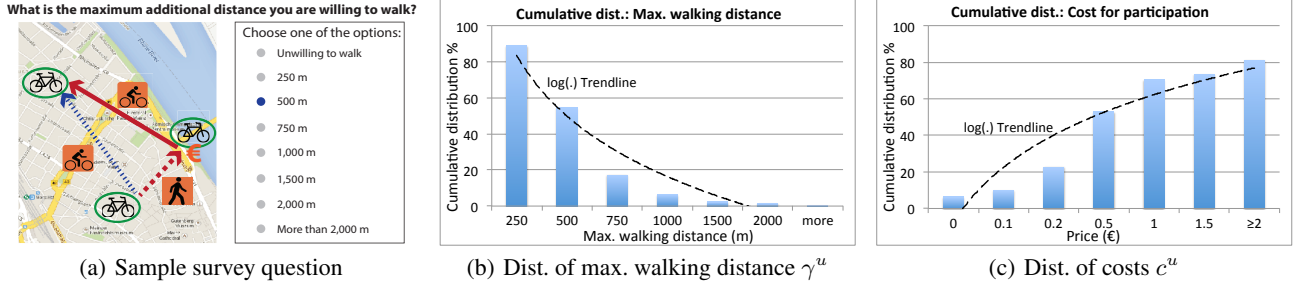


Figure 3: Survey study with customers of a real-world BSS¹ in a city of Europe where we deployed our incentive system

iteration n , and used to compute the optimal price p^n to offer. The mechanism receives the binary feedback y^n which is then used to update the estimates for price p^n .

Truthfulness of Incentives System

As the system users may act strategically, it is important to analyze the truthfulness of the IDS to understand the real-world performance of the proposed system. Each user has a private cost c^u and location l_x^u and we now discuss the truthfulness of our system w.r.t. these two dimensions of the users' private information. The mechanism DBP-UCB is based on the posted-price model, where it is in the user's best interest to truthfully accept the offer whenever the offered price is higher than the private cost c^u . Furthermore, given the large scale of the system, the interaction with a single user does not significantly affect the learning behavior of the mechanism, thus making it behave truthfully in real-world BSSs.

Let us now analyze the truthfulness of IDS w.r.t the location information. One crucial aspect here is that no incentives are offered when the user's default station s_x^u is already among the candidate stations. This notably improves the efficiency (by over 50% in our simulations) of the IDS by avoiding to pay out incentives for stations that the user would have taken anyways. We discuss the implications of this on the truthfulness of IDS below. If the system has no way to infer s_x^u , the user may indeed act strategically by misreporting to obtain an incentive that would have not been offered otherwise. In the real deployment of our system, instead of allowing the user to declare her current location (l_x^u), it is directly retrieved from the deployed mobile-app which provides localization features. Based on this location information, the user's default pickup station can directly be inferred. This technical solution does not directly apply to the bike drop-off scenario where users explicitly declare their target location. However the drop-off location could potentially be inferred based on user's historical trips, making the system incentive-compatible in real-world deployments.

Experimental Evaluation

In this section, we carry out extensive experiments to understand the practical performance of our system.

Datasets and Experimental Setup

First, we describe our datasets and the BSS simulator.

Historical BSS dataset. For our simulations, we used a historical dataset from Boston's Hubway, made publicly

available by the Boston Metropolitan Area Planning Council². The Hubway dataset contains data collected between 28th July, 2011 and 1st October, 2012 with rich information about 95 stations, 694 bicycles, 552,030 rentals, and snapshots of the status of the BSS at regular intervals.

Survey study. We did a survey study among the customers of a real-world BSS¹ in a city of Europe where we also deployed our system. Our goal is to validate the realism of our model, as well as to obtain realistic statistics about the distribution of users' personal costs c^u and of their maximum walking range γ^u . The participants were asked generic questions about their rental behaviors (e.g., purpose and duration of the trips) and then we introduced them to the proposed incentives system asking various questions to elicit their preferences about private cost and walking distance. Figure 3(a) illustrates a survey question, and the distributions obtained from the survey for γ^u and c^u are shown in Figures 3(b) and 3(c) respectively. Note that the survey also contained an additional choice of "unwilling to walk" or "unwilling to participate at any cost" that accounted for roughly 20% of the response for both questions.

Simulator. We built a complete simulator of a real BSS based on Hubway's historical data as well as the survey data from customers. Each simulation starts by taking a snapshot of the status of the BSS, i.e., the number of bikes at each station retrieved from the historical dataset and runs the simulation for a total of 30 days. The BSS simulator generates users' trip events by sampling from the distribution learnt from the historical data. Lastly, the simulator associates each rental to a user with cost c^u and γ^u sampled from the distributions obtained from the survey. For the TRUCKS, we used the myopic greedy policy as defined by Chemla, Meunier, and Pradeau (2013). Using the idea of *dynamic* (during rush hours) and *static* (during off hours) repositioning from Raviv and Kolka (2013), our policies operate in time from 12:00 p.m. and 3:00 p.m. and then from mid-night to morning. We assume that the trucks entail a fixed cost of 50 € per hour to the system and the policy allocates the number of hours based on the total budget.

Results

We now discuss the findings from our experiments.

Varying budget. We compare the performance of different policies by varying the allocated budget. We envision that the operators would deploy the incentives system in par-

²<http://hubwaydatachallenge.org/>.

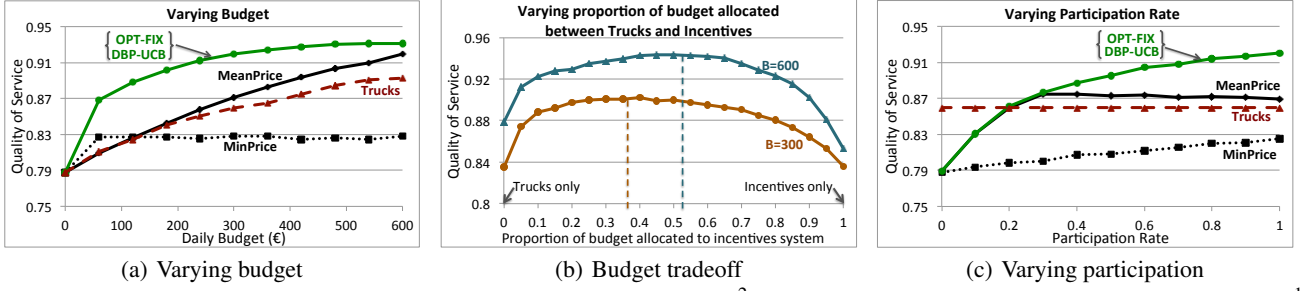


Figure 4: Simulation results based on historical data of Hubway BSS² and survey data from customers of real-world BSS¹

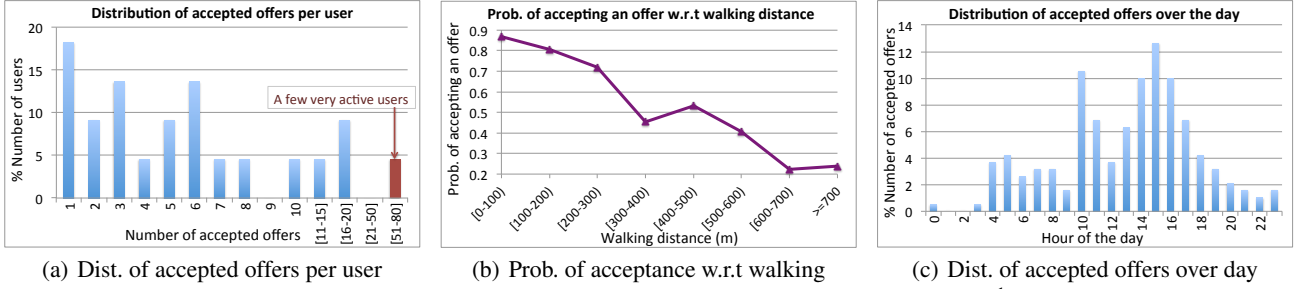


Figure 5: Results from deployment of our system through smartphone app on a real-world BSS¹ in beta-test phase for 30 days

allel to existing TRUCKS repositioning rather than as standalone policy. To take this into consideration, we compare the policies when deployed in addition to an already running TRUCKS policy allocated with 150 €, (i.e., equivalent of 3 hours of trucks that would run between 12:00 p.m. and 3:00 p.m.). Apart from comparing to TRUCKS, we also compare our proposed system with IDS running different baseline pricing mechanisms instead of DBP-UCB: i) MINIMUM mechanism selects the minimum price available c_{min} ; ii) MEAN mechanism offers the mean price obtained from the user's cost distribution. Our results in Figure 4(a) indicate that the proposed IDS with DBP-UCB pricing mechanism compares favorably w.r.t. TRUCKS, as well as other baseline pricing mechanisms. The performance of DBP-UCB almost matches OPT-FIX, as it converges very fast to the optimal price in less than a day of simulation time. Note that allocating more budget B to TRUCKS is simply equivalent to running policy with total of $B + 150$ € budget.

Budget tradeoff between IDS and TRUCKS. This experiment aims to understand how to trade off budget investment between TRUCKS and the incentives system. Figure 4(b) illustrates that the best service levels are achieved when the two repositioning services are run in complement to each other, and shows the high potential of adding our incentives system to the existing BSS infrastructure. In fact, these two policies have complementary effects. While the trucks tackle the imbalances at a macro level by moving the bicycles from a part of the city to the other, the effect of the incentives system is rather dynamic and at a micro level, incorporating the traffic flow and fluctuating demands.

Varying user participation. The previous experiments assume that all the BSS customers participate in the incentives system. However, this assumption may not hold in reality, especially in the testing phase of the system. With the same setup as used in Figure 4(a), we vary the participation

rate of users between 0% and 100%, with budget fixed at 300 €. Figure 4(c) illustrates that the proposed system can surpass standalone TRUCKS policy with 20% participation.

Deployment

We deployed our incentives system on a real-world BSS in a beta-test phase for a period of 30 days in a city of Europe, in collaboration with a large scale bike sharing company¹. We designed a smartphone app as an interface between the users and the incentives system. The system itself was integrated with the BSS infrastructure through their APIs. The participants were offered monetary incentives for the bike pickup scenario, with payments directly transferred to the customer accounts. Given the small pool of participants in the testing phase limiting the process of learning, we adopted a simple pricing mechanism based on OPT-FIX. Using the approximate users' cost distribution, denoted as \tilde{F} , obtained through the survey study, we computed the approximation of the optimal price \tilde{p}^{OPT} using Equation 1 by replacing F with \tilde{F} . We now present the findings from this deployment.

Participation. Figure 5(a) represents the histogram of the number of incentives obtained by each participant. The result shows that most participants collected five or less incentives each while few particularly active users collected more than 50 incentives each.

Reaction to incentives. In total, the acceptance rate of the incentive offers over all participants was about 60%. Here, the rejection of an offer corresponds to the case of a user rejecting an incentive to pick up a bike at an offered target station and starting a rental somewhere else. Furthermore, as we recorded the information about the user's location when an offer is made, we can compute the distance that the users were required to walk for a given offer. Figure 5(b) shows the probability of an offer being accepted as decreasing with respect to the distance to be walked and matches closely with the survey data in Figure 3(b).

Time and space distribution. Figure 5(c) shows that incentives were not collected equally throughout the day, but most incentives were collected during the afternoon and at late morning. In Figure 2, the diameter of each station is proportional to the number of incentives collected by users at that station. The map shows that the majority of incentives were paid out at stations in the city center.

Conclusions

There is much potential in intelligent and self-sustainable systems that incentivize and empower their users to actively engage in the system processes. In this paper, we presented the architecture of an incentives system to engage BSS users in the bike repositioning process. We worked in collaboration with a bike sharing company, designed and built a working system integrated with the BSS infrastructure. We investigated the incentive compatibility of our system, and also designed a dynamic pricing mechanism DBP-UCB to optimize the efficiency of the incentives system. We evaluated the proposed system through extensive simulations using historical and user survey data, and deployed it through a smartphone app for a period of 30 days.

Acknowledgments. We'd like to thank bike sharing operator MVGmeinRad for promoting the surveys and trial to their customers, and ElectricFeel Mobility Systems for providing access to their software platform and infrastructure for running the trial. This research is supported in part by SNSF grant 200021_137971 and Nano-Tera.ch program as part of the Opensense II project.

References

- [2011] Borgnat, P.; Abry, P.; Flandrin, P.; Céline, R.; Rouquier, J.-B.; and Fleury, E. 2011. Shared bicycles in a city: a signal processing and data analysis perspective. *Advances in Complex Systems* 14:415–438.
- [2011] Büttner, J.; Mlasowsky, H.; and Birkholz, T. 2011. Optimising Bike Sharing in European Cities - A Handbook.
- [2012] Cepolina, E. M., and Farina, A. 2012. Urban car sharing: An overview of relocation strategies. *WIT Transactions on the Built Environment* 419–431.
- [2013] Chemla, D.; Meunier, F.; and Pradeau, T. 2013. Self-service bike sharing systems: simulation, repositioning, pricing.
- [2013] Chemla, D.; Meunier, F.; and Wolfler Calvo, R. 2013. Bike sharing systems: Solving the static rebalancing problem. *Discrete Optimization* 10:120–146.
- [2013] Clemente, M.; Fanti, M. P.; Mangini, A. M.; and Ukovich, W. 2013. The vehicle relocation problem in car sharing systems: modeling and simulation in a petri net framework. In *Application and Theory of Petri Nets and Concurrency*. Springer. 250–269.
- [2012] Côme, E., and Oukhellou, L. 2012. Model-based count series clustering for Bike-sharing system usage mining, a case study with the Vélib' system of Paris. *Transportation Research-Part C Emerging Technologies* 22:88.
- [2012] Correia, G. H. d. A., and Antunes, A. P. 2012. Optimization approach to depot location and trip selection in one-way carsharing systems. *Transportation Research Part E: Logistics and Transportation Review* 48:233–247.
- [2012] Fricker, C., and Gast, N. 2012. Incentives and regulations in bike-sharing systems with stations of finite capacity. *arXiv preprint arXiv:1201.1178* 1–25.
- [2009] Froehlich, J.; Neumann, J.; and Oliver, N. 2009. Sensing and Predicting the Pulse of the City through Shared Bicycling. In *IJCAI*, 1420–1426.
- [2011] George, D. K., and Xia, C. H. 2011. Fleet-sizing and service availability for a vehicle rental system via closed queueing networks. *European Journal of Operational Research* 211:198–207.
- [2011] GoDCgo. 2011. Capital Bikeshare's Reverse Rider Rewards. <http://www.godcgo.com/reverse-rider-rewards.aspx>.
- [2014] Han, Y.; Côme, E.; and Oukhellou, L. 2014. Towards Bicycle Demand Prediction of Large-Scale Bicycle Sharing System. In *Transportation Research Board 93rd Annual Meeting*.
- [2010] Kaltenbrunner, A.; Meza, R.; Grivolla, J.; Codina, J.; and Banchs, R. 2010. Urban cycles and mobility patterns: Exploring and predicting trends in a bicycle-based public transport system. *Pervasive and Mobile Computing* 6:455–466.
- [2009] Kek, A. G. H.; Cheu, R. L.; Meng, Q.; and Fung, C. H. 2009. A decision support system for vehicle relocation operations in carsharing systems. *Transportation Research Part E: Logistics and Transportation Review* 45:149–158.
- [1985] Lai, T. L., and Robbins, H. 1985. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics* 6(1):4–22.
- [2014] Meddin, R., and DeMaio, P. 2014. The Bike-sharing World - First Week of June 2014. <http://bike-sharing.blogspot.ch/2014/06/the-bike-sharing-world-first-week-of.html>.
- [2011] Nair, R., and Miller-Hooks, E. 2011. Fleet Management for Vehicle Sharing Operations. *Transportation Science* 45:524–540.
- [2013] Nair, R.; Miller-Hooks, E.; Hampshire, R. C.; and Bušić, A. 2013. Large-Scale Vehicle Sharing Systems: Analysis of Vélib'. *International Journal of Sustainable Transportation* 7:85–106.
- [2014] Pfrommer, J.; Warrington, J.; Schildbach, G.; and Morari, M. 2014. Dynamic Vehicle Redistribution and Online Price Incentives in Shared Mobility Systems. *IEEE Intelligent Transportation Systems Magazine*.
- [2013] Rainer-Harbach, M.; Papazek, P.; Hu, B.; and Raidl, G. R. 2013. Balancing bicycle sharing systems: A variable neighborhood search approach. In *Evolutionary Computation in Combinatorial Optimization*. Springer. 121–132.
- [2013] Raviv, T., and Kolka, O. 2013. Optimal inventory management of a bike-sharing station. *IIE Transactions* 45(10):1077–1093.
- [2013] Sadilek, A.; Krumm, J.; and Horvitz, E. 2013. Crowdphysics: Planned and Opportunistic Crowdsourcing for Physical Tasks. *SEA* 536–545.
- [2013] Schuijbroek, J.; Hampshire, R.; and van Hoeve, W.-J. 2013. Inventory rebalancing and vehicle routing in bike sharing systems.
- [2010] Singer, Y. 2010. Budget feasible mechanisms. In *FOCS*.
- [2013] Singla, A., and Krause, A. 2013. Truthful incentives in crowdsourcing tasks using regret minimization mechanisms. In *WWW*.
- [2014] Vélib'. 2014. How it works: Vélib'. <http://en.velib.paris.fr/How-it-works/Stations>.
- [2013] Waserhole, A.; Jost, V.; and Brauner, N. 2013. Pricing techniques for self regulation in Vehicle Sharing Systems. *Electronic Notes in Discrete Mathematics* 41:149–156.