

On the Bike Spreading Problem

Elia Costa

Dept. of Information Engineering, University of Padova, Italy

Francesco Silvestri¹   

Dept. of Information Engineering, University of Padova, Italy

Abstract

A free-floating bike-sharing system (FFBSS) is a dockless rental system where an individual can borrow a bike and returns it anywhere, within the service area. To improve the rental service, available bikes should be distributed over the entire service area: a customer leaving from any position is then more likely to find a near bike and then to use the service. Moreover, spreading bikes among the entire service area increases urban spatial equity since the benefits of FFBSS are not a prerogative of just a few zones. For guaranteeing such distribution, the FFBSS operator can use vans to manually relocate bikes, but it incurs high economic and environmental costs. We propose a novel approach that exploits the existing bike flows generated by customers to distribute bikes. More specifically, by envisioning the problem as an Influence Maximization problem, we show that it is possible to position batches of bikes on a small number of zones, and then the daily use of FFBSS will efficiently spread these bikes on a large area. We show that detecting these zones is NP-complete, but there exists a simple and efficient $1 - 1/e$ approximation algorithm; our approach is then evaluated on a dataset of rides from the free-floating bike-sharing system of the city of Padova.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis; Theory of computation → Approximation algorithms analysis; Information systems → Data mining

Keywords and phrases Mobility data, bike sharing, bike relocation, influence maximization, NP-completeness, approximation algorithm

Digital Object Identifier 10.4230/OASICS.ATMOS.2021.5

Related Version Updates to the current paper will be made available on arXiv.

arXiv link: <https://arxiv.org/abs/2107.00761>

Supplementary Material The source code and input graphs used for the experiments are publicly available.

Software (Source Code): <https://github.com/AlgoUniPD/BikeSpreadingProblem>

archived at `swh:1:dir:d112339ca8a7c5bdad285cbb441537fe32f58734`

Funding The paper was partially supported by UniPD SID18 grant, PRIN17 20174LF3T8, MIUR “Departments of Excellence”.

Acknowledgements The authors would like to thank the Municipality of Padova for providing us the dataset with rides of the free-floating service in Padova. We are also grateful to Pietro Rampazzo for providing us the grids of the city of Padova and useful suggestions on plotting maps.

1 Introduction

A *bike-sharing system* (BSS) is a service where an individual can rent a bike and return it after a short term. Nowadays, almost all large cities have adopted a BSS as it is a sustainable transportation system that helps improving air pollution, public health, and traffic congestion [20]. The first BSS dates back to the 1960s (with Witte Fietsen in Amsterdam), and we have seen an explosion of BSSs in the last decade, with now about 2000

¹ Corresponding author.



5:2 On the Bike Spreading Problem

operators currently managing more than 9.7 million bikes around the world [13]. The majority of BSSs are now connected to IT systems that allow to borrow bikes from a smartphone and to collect data on users and rides. Similar systems exist for e-bikes, scooters, mopeds, and cars.

There are two major approaches to BSSs: station-based and free-floating systems. A *station-based bike-sharing system* (SBBSS) represents the most common approach, where a user borrows a bike from a dock and returns it at another dock belonging to the same system; in an SBBSS, the set of origins and destinations of all rides is small and coincides with dock positions. On the other hand, a *free-floating bike-sharing system* (FFBSS) is a sharing model with no docks: each bike has an integrated lock that can be opened on demand with a smartphone app and returned just by closing the lock. As there are no more fixed docks, bikes can be positioned anywhere (hopefully, respecting traffic codes) and the origin/destination of a ride can be any position within the service area. FFBSS is also an interesting solution for the first and last kilometer problem in multimodal transportation cities [16] since the average walking distance of a user to the closest free-floating bike is shorter than SBBSS.

The distribution of bikes in the service area is crucial for to increase user satisfaction. In SBBSS, the main challenge is dealing with hotspots, specifically zones where several rides start (sources) or end (sinks), such as train stations or university campuses. Hotspots are critical: sources and sinks might suffer, respectively, from the lack of available bikes and parking slots in a deck. The operator needs to detect and manage these hotspots: bikes should be collected from sinks and repositioned on sources. The number of hotspots is usually small and several efficient computational strategies have been investigated (see e.g. references in [16]). In FFBSS, hotspots are still critical although customers do not experience the lack of empty slots in the returning docks as in SBBSS.

In addition to dealing with hotspots, FFBSS needs also to guarantee that the entire service area is covered by bikes so that a user leaving from any point can find a near available bike. A study [7] has indeed shown that every additional meter of walking to a shared bike decreases a user's likelihood of using a bike by 0.194% for short distances ($\leq 300m$) and 1.307% for long distances ($> 300m$), implying that a user walking a distance $> 500m$ for reaching the closest bike is highly unlikely to use the system. Moreover, if bikes are well distributed all over the service area, the spatial equity improves since the benefits of the service are not a prerogative of just some zones (e.g., city center) [14]. Assume the service area to be split into zones (e.g., quadrants) where the diameter of each zone is considered a reasonable walking distance (e.g., $\leq 500m$): then, the desired goal is that each zone has a sufficient number of bikes.

To distribute bikes over the service area, an FFBSS operator could employ a fleet of vans to manually position bikes in each zone: however, this solution might be economically and environmentally unfeasible due to a large number of zones. In this paper, we provide a novel and alternative approach that distributes bikes over the service area by exploiting the existing customers' bike flows and hence reducing intervention by the FFBSS operator. The idea is to detect a small number of zones, named *seeds*, where bikes are more likely to be spread over the service area by the regular activity of customers during a given time interval. The seeds represent the positions where the FFBSS operator can position batches of bikes, which will then be spread over the entire area by customers, without further interventions from the operator. Formally, we modeled this approach as a variant of the Influence Maximization (IM) problem, which we name the *Bike Spreading (BS) problem*.

To detect these seeds, we first define a weighted graph representing mobility flows: nodes are zones, and weighted edges represent the probability that a bike moves from one node to another. Then, we introduce a diffusion model to analyze how bikes move and a spread score to evaluate the quality of the final bike distribution. Finally, we detect a small subset of nodes that maximizes the spread score and position bikes on these nodes. This formulation results in an NP-complete problem, but we show that there exists a $1 - 1/e$ -approximation algorithm to the problem thanks to some properties (e.g., submodularity and monotonicity).

More specifically, the results provided in the paper are the following:

- In Section 3, we introduce the Bike Spreading problem and formalize two versions called T-BS and U-BS: the T-BS version aims at maximizing the number of zones with a minimum amount $\gamma > 0$ of bikes, while the goal of U-BS is to uniformly distribute bikes in the service area.
- In Section 4, we analyze the theoretical properties of T-BS and U-BS: we show their NP-completeness and that U-BS satisfies the monotonicity and submodularity properties. By these properties and the result by Nemhauser et al. [15], we get a simple greedy algorithm providing a $1 - 1/e$ -approximation for U-BS.
- In Section 5, we experimentally investigate the BS problem by using data from the free-floating bike service of the city of Padova (Italy). We analyze the performance and quality of the solution of the greedy algorithm, compare the U-BS and T-BS problems, and make some empirical considerations on the BS problem.

2 Preliminaries

2.1 Free-floating bike-sharing service

The studies related to bike-sharing systems mainly involve two topics: demand prediction and rebalancing. Demand analysis involves the understanding of user behavior and providing the most appropriate service (e.g., [12]). Rebalancing has mainly focused on station-based systems (see e.g. [11, 5, 19, 16]), while only a few works have addressed free-floating bikes. Reiss and Bogenberger [17] investigated the relocation strategy and a validation method on Munich's FFBS; their approach focused on finding the best zones where to relocate bikes to satisfy user demand and minimize bikes' idle time. Pal and Zang [16] and Usamaa et al. [18] focused on finding the best route of relocation vans under costs and time constraints; in [18], picking up faulty bikes was also included. Finally, Caggiani et al. [3] proposed a forecast model aiming at reducing the number of times when a zone has fewer bikes than necessary. These works focused on computing an optimal route to collect and relocate bikes, and on detecting hotspots where to relocate more bikes. To the best of our knowledge, no previous works have studied how to spread bikes over the entire service area and exploited the mobility graph as in our work.

2.2 The Influence Maximization problem

The *Influence Maximization* (IM) problem [8] is widely used in social network analysis to detect the most influential users that can efficiently spread information, like news or advertisements; IM is also used in epidemiology for analyzing how infections evolve via human interactions. However, differently from news and infections, bikes do not replicate: we then need to redefine the IM framework to deal with a fixed amount of objects spreading over the graph. In general, an Influence Maximization problem consists of three components:

5:4 On the Bike Spreading Problem

- A directed and weighted graph where vertexes represent users, and edges represent the paths where information can propagate from a given node.
- A diffusion model that describes how information moves in the graph. The model specifies the initial status (e.g., the nodes that initially contain the information) and how information distributes. The model advances in steps: in each step, a node detects which other nodes in the neighborhood will receive the information and propagates it.
- An evaluation function $\sigma(\Upsilon)$, which receives a description Υ of how the information is distributed in the graph and provides a non-negative real value. Large values denote better and more desired distributions; e.g., $\sigma(\Upsilon)$ can represent the total number of nodes that have seen some information.

The goal of IM is to find an initial distribution that maximizes the evaluation function $\sigma(\Upsilon)$ after a given number of steps.

The diffusion model is the critical part of IM problems, indeed they delineate the way nodes influence each other. The most commonly used diffusion models [4, 10] are Independent Cascade, Linear Thresholds, Triggering, and Time Aware: these model are progressive models, in the sense that once a node has been influenced it cannot change its status. There are also non-progressive diffusion models: some examples are the Susceptible-Infected-Susceptible (SIS) models [9], widely used in epidemiology.

The bike spreading problem can be viewed as a IM problem where the diffusion model propagates objects (i.e., bikes). The main difference is that objects cannot be replicated and their amount is constant. In contrast, previous works on information or infection diffusion allow replicas: for instance, an infected person can infect many other persons and an image can be shared with friends in a social network.

2.3 Approximating submodular functions

Nemhauser et al. [15] proved that a non-negative, monotone submodular function can be efficiently approximated by a simple greedy algorithm, within a factor $1 - 1/e \sim 0.63$. This result is used by Kempe et al. [8] to obtain an approximate solution for the IM problem under both Independent Cascade and Linear Threshold models.

Consider a function f mapping a set of elements in U to a non-negative real value, i.e. $f : U^* \rightarrow \mathbb{R}^+$. Intuitively, a function f is monotone if, by expanding a given input set, the value of the function does not decrease; a function f is submodular if the marginal gain does not increase when adding more elements to the input set.

► **Definition 1** (Monotonicity). *Given a function $f : U^* \rightarrow \mathbb{R}^+$ where U is a set of elements, function f is said to be monotone if, for every $S \subseteq U$ and $v \in U$, we have $f(S \cup \{v\}) \geq f(S)$.*

► **Definition 2** (Submodularity). *Given a function $f : U^* \rightarrow \mathbb{R}^+$ where U is a set of elements, function f is said to be submodular if, for every $S, T \subseteq U$ with $S \subseteq T$ and $v \in U \setminus T$, we have $f(S \cup \{v\}) - f(S) \geq f(T \cup \{v\}) - f(T)$.*

The aforementioned result by Nemhauser et al. [15] is the following:

► **Theorem 3** ([15]). *Let $f : U^* \rightarrow \mathbb{R}^+$ be a non-negative, monotone submodular function where U is a set of elements. Let $k \geq 1$ be a given value and S^* be a set of k elements in U maximizing the value of f among all sets of size k . Then, there exists a greedy algorithm that returns a set S of k elements such that $f(S) \geq 1 - (1 - 1/k)^k f(S^*) \geq (1 - 1/e)f(S^*)$, where e is the base of the natural logarithm.*

The greedy algorithm is quite simple and it consists of extending a given set S with the node u^* that maximizes the marginal gain, i.e., $u^* = \arg \max_{u \in U} \{f(S \cup \{u\}) - f(S)\}$ until we get a set of k elements.

```

greedy( $f, k$ )
 $S = \emptyset$ 
for  $i = 0$  to  $k - 1$ 
     $u^* = \arg \max_{u \in U} \{f(S \cup \{u\}) - f(S)\}$ 
     $S = S \cup \{u^*\}$ 
return  $S$ 

```

The algorithm takes time $O(k\Delta)$ where Δ is the maximum cost of finding the element of U that maximizes the marginal gain.

3 The Bike Spreading problem

In this section we formalize the Bike Spreading problem. We first propose a general definition based on the Influence Maximization problem, and then we consider two special cases, named U-BS and T-BS problems, that target specific bike distributions.

The Bike Spreading (BS) problem can be viewed as a special case of the Influence Maximization problem, where the entity that is distributed across the graph are items that cannot be replicated, differently from news and infections.² We represent a city as a *directed* graph $G = (V, E)$: V is the set of nodes, where every node represents a different zone of the city; E is the set of directed weighted edges capturing the probability of moving from one zone to another. We define $n = |V|$ and $m = |E|$, and let $\Gamma^{IN}(v) = \{w \in V | (w, v) \in E\}$ and $\Gamma^{OUT}(v) = \{w \in V | (v, w) \in E\}$ denote the set of nodes for which v is the destination or the source, respectively. The weight of an edge $e = (u, v) \in E$ represents the *probability* p_e that a bike in node u moves from u to v : intuitively, p_e is the probability that an user renting a bike in zone u ends her/his ride into zone v . We use self-loops to represent bikes that stay in the same node (i.e., for bikes that are not used or are borrowed for a ride starting and ending in the same node). Then, for each node $u \in V$, we have:

$$\sum_{v \in \Gamma^{OUT}(u)} p_{(u,v)} = 1. \quad (1)$$

We assume bikes to be initially positioned on $k \geq 1$ nodes in groups of $L \geq 1$ bikes, and we refer to the initial set of these k nodes as (k, L) -*seed*. Bikes can be damaged or stolen in a free-floating bike system, however, this does not significantly affect the total number of bikes in the system: therefore, we assume the total number of bikes in the graph to be fixed, and we let B denote the total number of bikes, i.e. $B = k \cdot L$.

The diffusion model unfolds in $\tau \geq 1$ discrete steps. At any step $1 \leq t \leq \tau$ and for each vertex $v \in V$, we define the *load* $\ell_{k,L}(v, t, \mathcal{S}) \geq 0$ which represents the (expected) number of bikes in node v after t steps starting from a (k, L) -seed set \mathcal{S} . We let $\ell_{k,L}(v, 0, \mathcal{S})$ denote the initial loads: $\ell_{k,L}(v, 0, \mathcal{S}) = L$ for each $v \in \mathcal{S}$, and $\ell_{k,L}(v, 0, \mathcal{S}) = 0$ otherwise. At any time $0 < t \leq \tau$, bikes move according to edge directions and probabilities; each vertex load $\ell_{k,L}(v, t, \mathcal{S})$ is updated as follows:

$$\ell_{k,L}(v, t, \mathcal{S}) = \ell_{k,L}(v, t-1, \mathcal{S}) + \Delta^{IN}(\mathcal{S}) - \Delta^{OUT}(\mathcal{S})$$

² The model presented in our paper is not atomic: we allow a bike to be “split” since the value of a node should be understood as the expected number of bikes in that node.

5:6 On the Bike Spreading Problem

where the two rightmost terms are defined as

$$\Delta^{IN}(\mathcal{S}) = \sum_{u \in \Gamma^{IN}(v)} p_{(u,v)} \ell_{k,L}(u, t-1, \mathcal{S}), \quad \Delta^{OUT}(\mathcal{S}) = \sum_{u \in \Gamma^{OUT}(v)} p_{(v,u)} \ell_{k,L}(v, t-1, \mathcal{S}).$$

Intuitively, the bikes in a node v are partitioned among the outgoing edges according to their probabilities, and then the load of v is decreased by the number of bikes leaving the node (i.e., $\Delta^{OUT}(\mathcal{S})$) and increased by the number of bikes entering in the node (i.e., $\Delta^{IN}(\mathcal{S})$). As $\sum_{u \in \Gamma^{OUT}(v)} p_{(v,u)} = 1$ due to self-loops, the above update can be rewritten as follows:

$$\ell_{k,L}(v, t, \mathcal{S}) = \sum_{u \in \Gamma^{IN}(v)} p_{(u,v)} \ell_{k,L}(u, t-1, \mathcal{S}).$$

We let $\mathcal{L}_{k,L}(t, \mathcal{S}) = (\ell_{k,L}(v_0, t, \mathcal{S}), \dots, \ell_{k,L}(v_{n-1}, t, \mathcal{S}))$ denote the loads of all nodes in G after $1 \leq t \leq \tau$ steps and with the (k, L) -seed \mathcal{S} . We remark that $\ell_{k,L}(v, t, \mathcal{S})$ represents the expected number of bikes in node v after t steps, where all bikes in a node are spread among the neighbors nodes in every time step uniformly at random according to edge probabilities.

To measure the quality of bike distribution among nodes, we introduce the notion of spread. The *spread* $\sigma(\mathcal{L}_{k,L}(t, \mathcal{S}))$ after t steps and with (k, L) -seed \mathcal{S} is a function $\mathbb{R}^n \rightarrow \mathbb{R}$ that evaluates the quality of loads $\mathcal{L}_{k,L}(t, \mathcal{S})$. The actual formulation of the spread function depends on the desired distribution in a free-floating bike system. We will see two examples of spread aiming at maximizing the number of nodes with a given minimum load, and at uniformly distributing bikes among all nodes of the graph. We are now ready to define the Bike Spreading problem as:

► **Definition 4.** *Given a directed graph $G = (V, E)$, positive integers $L \geq 1$, $k \geq 1$ and $\tau \geq 1$, and spread $\sigma(\cdot)$, the Bike Spreading (BS) problem asks for the (k, L) -seed \mathcal{S}^* that maximizes $\sigma(\mathcal{L}_{k,L}(\tau, \mathcal{S}))$. More specifically, the mathematical formulation of the problem is*

$$\mathcal{S}^* = \arg \max_{\mathcal{S} \subseteq V, |\mathcal{S}|=k} \sigma(\mathcal{L}_{k,L}(\tau, \mathcal{S}))$$

such that:

- $\ell_{k,L}(v, 0, \mathcal{S}) = L, \quad \forall v \in \mathcal{S}$
- $\ell_{k,L}(v, 0, \mathcal{S}) = 0, \quad \forall v \in V \setminus \mathcal{S}$
- $\ell_{k,L}(v, t, \mathcal{S}) = \sum_{u \in \Gamma^{IN}(v)} p_{(u,v)} \ell_{k,L}(u, t-1, \mathcal{S}), \quad \forall v \in V \text{ and } t = 1, \dots, \tau$

We remark that in this paper we focus on small values of τ (i.e., the maximum number of steps) and thus we do not analyze the convergence of the spread for $\tau \rightarrow +\infty$. As the distribution should happen within a couple of hours from the positioning of bike batches, we expect each bike to be used just a few times and thus $\tau \leq 5$ from a practical point of view.

We now describe the two spread functions used in the paper, namely T-BS and U-BS.

T-BS version

The first example of spread leverages on the idea that a zone is considered well served by the free-floating bike system if there are at least γ bikes (in expectation), where $\gamma > 0$ is a given threshold value. Therefore, the spread counts the number of nodes (i.e., zones) that have at least γ bikes:

$$\sigma_{\gamma}^{(T)}(\mathcal{L}_{k,L}(\tau, \mathcal{S})) = |\{v \in V \mid \ell_{k,L}(v, \tau, \mathcal{S}) \geq \gamma\}| \quad (2)$$

We refer to the BS formulation with the spread in Equation 2 as *Threshold BS (T-BS)* problem.

U-BS version

The T-BS problem is an all-or-nothing approach, where only nodes receiving a sufficiently large number of bikes are relevant. However such a solution would not penalize skewed distributions: for instance, T-BS gives the same score to a distribution of n nodes with load 2γ and a skewed distribution of $n - 1$ nodes with load γ and one node with load $(n + 1)\gamma$. Therefore, if the goal is to maximize the area covered by the service, the ideal distribution should be the uniform distribution with B/n bikes per node. We thus introduce the following spread:

$$\sigma^{(U)}(\mathcal{L}_{k,L}(\tau, \mathcal{S})) = \sum_{v \in V} \sqrt{\ell_{k,L}(v, \tau)} \quad (3)$$

The maximum value of $\sigma^{(U)}(\mathcal{L}_{k,L}(\tau, \mathcal{S}))$ is reached when bikes are equally distributed over the graph, that is $\ell_{k,L}(v, \tau) = B/n$. We refer to the BS formulation with the spread in Equation 3 as *Uniform BS (U-BS)* problem.

4 Theoretical analysis

In this section, we first prove the NP-completeness of T-BS and U-BS problems. Then we show that U-BS satisfies the monotonicity and submodularity properties and hence, by the result of Nemhauser et al. [15], there exists a greedy algorithm providing a $(1 - 1/e)$ -approximate solution to U-BS.

4.1 NP-completeness

T-BS is NP-complete

We prove the NP-completeness of T-BS with a reduction from the *Minimum Dominating Set (MDS)* problem which is known to be NP-complete even in graphs with constant vertex degree $d \geq 3$ [2, 1]. Let $G = (V, E)$ be a simple, connected, undirected graph with degree $d \geq 3$ for all nodes in V . A subset $S \subseteq V$ is a *dominating set* if for every vertex $v \in V \setminus S$, there exists a vertex $u \in S$ such that $(v, u) \in E$; that is, every vertex outside S has at least one neighbor in S . A minimum domination set of G is a domination set of the smallest possible size, and we refer to its size with *domination number* $\gamma(G)$. Since the degree of every node in G equals d , we have $\gamma(G) \leq n - d + 1$, because in a set $S \subseteq V$, $|S| \geq n - d + 1$, every node has at least one neighbor in S .

The decision problem associated with MDS is defined as:

- INSTANCE: a undirected graph $G = (V, E)$ with fixed degree $d \geq 3$ and an integer k with $0 < k < n - d + 1$.
- QUESTION: Does a dominating set with $\gamma(G) \leq k$ exist?

The T-BS problem can be represented by the following decision problem:

- INSTANCE: a directed and weighted graph $G = (V, E)$ that satisfies Equation 1, integers k, L, γ, τ and λ with $0 < k < n, L \geq 1, \gamma > 0, \tau \geq 1$, and $0 < \lambda \leq n$.
- QUESTION: Does a (k, L) -seed $\mathcal{S} \subseteq V$ exist in G such that $\sigma_{\gamma}^{(T)}(\mathcal{L}_{k,L}(\tau, \mathcal{S})) \geq \lambda$?

The reduction of MDS to T-BS is the following. Given an instance of the MDS problem on a graph G with degree $d \geq 3$ and integer k , we construct a new weighted and directed graph G' for the T-BS problem by directing all edges of G in both directions and by adding a self-loop to each vertex, obtaining a new graph $G' = (V, E')$ with outdegree $d' = d + 1$.

The probability of each edge is then set to $p_e = 1/d'$. We then solve the T-BS problem on G' with a seed with size k and $L = d'$ bikes in each vertex in the seed, threshold $\gamma = 1$, step number $\tau = 1$, and $\lambda = n$. The reduction returns yes to the d -MDS problem if and only if T-BS returns yes, that is if there exists a (k, d') -seed set \mathcal{S} in G' with $\sigma_1^{(T)}(\mathcal{L}_{k,d'}(1, \mathcal{S})) \geq n$.

► **Theorem 5.** *The T-BS problem is NP-complete.*

Proof. The T-BS problem is in NP since a solution can be verified in polynomial time as follows. Let A be the adjacency matrix of a graph G , then matrix A^τ represents the percentage of bikes in a node u that reach a node v , for each $u, v \in V$, using paths of length τ (including possibly self-loops). Matrix A^τ can be computed in $O(n^3 \log \tau)$ time with the doubling trick. Given a seed \mathcal{S} and an n -dimensional vector $\ell_{\mathcal{S}}$ encoding the initial loads, then the loads at step τ can be computed with $A^\tau \cdot \ell_{\mathcal{S}}$ in $O(n^2)$ time. Therefore $\sigma_\gamma^{(T)}(\mathcal{L}_{k,L}(\tau, \mathcal{S}))$ for a given seed can be computed in $O(n^3 \log \tau)$ time.

We now show the correctness of the reduction. We first prove that, if there exists a dominating set \mathcal{S} of size $k' \leq k$ in G , then there exists a (k, d') -seed in G' giving $\sigma_1^{(T)}(\mathcal{L}_{k,d'}(1, \mathcal{S})) \geq n$. Let us assume for simplicity that \mathcal{S} has size $k' = k$: it suffices to add nodes to \mathcal{S} till reaching size k . Since there are $L = d'$ bikes in each node of the seed set, $p_e = 1/d'$ and there are exactly d' outgoing edges from each node, we have that each outgoing edge of a node in \mathcal{S} is crossed by one bike. By definition, the dominating set \mathcal{S} covers each vertex in $V \setminus \mathcal{S}$ with at least one edge; moreover, each node in \mathcal{S} has a self-loop. Therefore, each node in V receives at least one bike and hence $\sigma_1^{(T)}(\mathcal{L}_{k,d'}(1, \mathcal{S})) = n$.

Conversely, if a dominating set \mathcal{S} of size at most k does not exist in G , then there cannot be a set \mathcal{S}' of size k in G' with $\sigma_1^{(T)}(\mathcal{L}_{k,d'}(1, \mathcal{S}')) \geq n$. Indeed, since a dominator set of size $\leq k$ does not exist, it means that for any set \mathcal{S}' of $\leq k$ nodes from V there is at least one node $v \in V$ which is not adjacent to nodes in \mathcal{S}' . Therefore, for any seed set \mathcal{S}' of k nodes there exists a node v not receiving any bike; it follows that $\ell_{k,d'}(v, 1, \mathcal{S}') = 0$ and thus $\sigma_1^{(T)}(\mathcal{L}_{k,d'}(1, \mathcal{S}')) \leq n - 1$. ◀

U-BS is NP-complete

We use a reduction from the *Exact Cover by 3-Sets (X3C)* problem, as MDS does not work for U-BS due to its spread function. X3C is NP-complete [6], and consists of a covering problem with sets of three elements. Given a set X with $|X| = 3q$, for some integer $q \geq 1$, and a collection C of 3-element subsets of X , X3C requires to decide if there exists a subset $C' \subseteq C$ such that C' covers X and every element of X occurs in exactly one set in C' (i.e., C' is an *exact cover* of X). For clarity, consider the following example: let $X = \{1, 2, 3, 4, 5, 6\}$ and $C = \{\{1, 2, 3\}, \{2, 3, 4\}, \{1, 2, 5\}, \{2, 5, 6\}, \{1, 5, 6\}\}$; then, the collection $C' = \{\{2, 3, 4\}, \{1, 5, 6\}\} \subset C$ is an exact cover because each element in X appears exactly once. Note that, if $C = \{\{1, 2, 3\}, \{2, 4, 5\}, \{2, 5, 6\}\}$, then any collection C' cannot be an exact cover: indeed, every pair of sets in C shares at least one entry, and hence every collection C' covers at least one element twice or more. Note that if we do have an exact cover, C' will contain exactly q elements.

The decision problem associated with X3C is:

- INSTANCE: a set X , with $|X| = 3q$ for some integer $q \geq 1$, a collection C of 3-element subsets of X .
- QUESTION: does a set $C' \subseteq C$ exist such that every element of X occurs in exactly one member of C' ?

The decision problem for U-BS is:

- **INSTANCE:** a directed and weighted graph $G = (V, E)$ that satisfies Equation 1, integers k, L, τ, λ with $0 < k < n, L \geq 1, \tau \geq 1, \lambda \geq 0$.
- **QUESTION:** Does a (k, L) -seed $\mathcal{S} \subseteq V$ exist in G such that $\sigma^{(U)}(\mathcal{L}_{k,L}(\tau, \mathcal{S})) \geq \lambda$?

The reduction from X3C to U-BS is the following. Given an instance of the X3C, defined by a set of $3q$ elements $X = \{x_1, \dots, x_{3q}\}$ and a collection $C = \{c_1, \dots, c_r\}$ of 3-element subsets of X , we build a directed and weighted graph $G = (V, E)$ for the U-BS problem as follows. With a slight abuse of notation, we let $V = C \cup X$, that is, each element in X and each set in C are represented by a node in V . Then we set $E = E_1 \cup E_2$: E_1 contains an edge (c_i, x_j) if $x_j \in c_i$, for each $c_i \in C$ and $x_j \in X$; E_2 contains a self-loop for each node $x_i \in X$. We observe that nodes in X have only one outgoing edge, while nodes in C have three outgoing edges: then, we set the probability of each self-loop in E_2 to 1, while we set $p_e = 1/3$ for all the remaining edges in E_1 . We then run the U-BS problem on G with $k = q$ nodes in the seed, $L = 3$ bikes per node in the seed, $\tau = 1$ steps, and $\lambda = 3q$, and we answer yes to X3C if and only if U-BS returns yes, that is if there exists a $(q, 3)$ -seed set \mathcal{S} in G with $\sigma^{(U)}(\mathcal{L}_{q,3}(1, \mathcal{S})) \geq 3q$.

► **Theorem 6.** *The U-BS problem is NP-complete.*

Proof. U-BS is in NP as a solution can be checked in polynomial time $O(n^3 \log \tau)$ as shown in the proof of Theorem 5.

We now prove the correctness of the reduction. We first observe that any exact cover for X3C must contain q entries from C , otherwise X is not covered or an element in X is covered by more than one set in C' . We now prove that, if the X3C problem contains an exact cover C' , then there exists a $(q, 3)$ -seed \mathcal{S} in G with $\sigma^{(U)}(\mathcal{L}_{q,3}(1, \mathcal{S})) \geq 3q$. Let \mathcal{S} be a seed set given by the nodes c_i representing sets in C' . As each node $x \in X$ is covered by exactly one node in C' and the outgoing degree of a node in C' is 3, we have that x receives one bike after the first step. Then: $\sigma^{(U)}(\mathcal{L}_{q,3}(1, \mathcal{S})) = \sum_{v \in V} \sqrt{\ell_{q,3}(v, 1, \mathcal{S})} = \sum_{v \in C} 0 + \sum_{v \in X} 1 = 3q$.

Assume now that X3C does not have an exact cover: then any C' of q sets from C does not cover at least one point of X and covers at least one point of X more than once. Assume by contradiction that U-BS returns a seed \mathcal{S} of $k = q$ nodes with $\sigma^{(U)}(\mathcal{L}_{q,3}(1, \mathcal{S})) \geq 3q$. We claim that \mathcal{S} has no nodes in X . If \mathcal{S} has a node $x \in X$, then the load of x after one step is $h = \ell_{q,3}(x, 1, \mathcal{S}) = \sqrt{3}$ since every node in X has only the self-loop as outgoing edge. Since bikes positioned in nodes of C' move in nodes of X after one step, we get:

$$\sigma^{(U)}(\mathcal{L}_{q,3}(1, \mathcal{S})) = \sqrt{h} + \sum_{v \in X \setminus \{x\}} \ell_{q,3}(v, 1, \mathcal{S}).$$

By the concavity of the square root, the right summation is maximized when all loads are equal and, since $|X| = 3q$, we get $\sigma^{(U)}(\mathcal{L}_{q,3}(1, \mathcal{S})) \leq \sqrt{h} + \sqrt{(3q-1)(3q-h)}$. The right term of the inequality decreases for $h > 1$, and then $\sigma^{(U)}(\mathcal{L}_{q,3}(1, \mathcal{S})) < 3q$, which is a contradiction. Therefore, we must have that \mathcal{S} contains only nodes in C . However, since there is no exact cover, at least one node in X must receive two or more bikes: by mimic the previous argument, we get that $\sigma^{(U)}(\mathcal{L}_{q,3}(1, \mathcal{S})) < 3q$. Therefore there is no $(q, 3)$ -seed giving $\sigma^{(U)}(\mathcal{L}_{q,3}(1, \mathcal{S})) \geq 3q$. ◀

4.2 Approximation algorithms

By the previous hardness results, we do not expect polynomial-time exact algorithms for the T-BS and U-BS problems. In this section, by showing that U-BS satisfies the monotonicity and submodularity properties, we get that the greedy solution in Section 2.3 gives a $(1 - 1/e)$ -approximation algorithm for U-BS. The T-BS version does not satisfy the submodularity

5:10 On the Bike Spreading Problem

property and thus similar theoretical guarantees cannot be provided: however, in the following section, we show that the greedy algorithm experimentally provides a good approximation even for T-BS.

We start with a technical lemma and then show that U-BS satisfies the monotonicity and submodularity properties.

► **Lemma 7.** *Let S and T be two seed sets with $S \subseteq T \subseteq V$, then $\ell_{k,L}(v, t, T) \geq \ell_{k,L}(v, t, S)$ for each $v \in V$ and $t \geq 0$.*

Proof. The proof follows by induction over the number of steps t . It is true in the base case when $\tau = 0$ since:

1. $\ell_{k,L}(v, 0, T) = \ell_{k,L}(v, 0, S) = 0$ for each $v \notin T$;
2. $\ell_{k,L}(v, 0, T) = \ell_{k,L}(v, 0, S) = L$ for each $v \in S$;
3. $\ell_{k,L}(v, 0, T) = L$ and $\ell_{k,L}(v, 0, S) = 0$ for each $v \in T \setminus S$.

Now suppose that the claim is true for $t \geq 0$. Then at step $t + 1$, we have:

$$\begin{aligned} \ell_{k,L}(v, t + 1, T) &= \sum_{w \in \Gamma^{tN}(v)} p_{(w,v)} \ell_{k,L}(w, t, T) \\ &\geq \sum_{w \in \Gamma^{tN}(v)} p_{(w,v)} \ell_{k,L}(w, t, S) = \ell_{k,L}(v, t + 1, S). \end{aligned}$$

The claim follows. ◀

► **Lemma 8.** *Given the U-BS problem with a seed set size k and given parameters τ , L independent of k , then $\sigma^{(U)}(\mathcal{L}_{k,L}(\tau, \mathcal{S}))$ is monotone and submodular.*

Proof. The monotonicity follows from Lemma 7 and the monotonicity of square root. We now consider submodularity. Since the parameters τ , L are given and are independent of k , we define $\sigma(\mathcal{S}) = \sigma^{(U)}(\mathcal{L}_{k,L}(\tau, \mathcal{S}))$ for notational simplicity. By Definition 2, we have to prove that $\sigma(\mathcal{S} \cup \{v\}) - \sigma(\mathcal{S}) \geq \sigma(T \cup \{v\}) - \sigma(T)$ for all $v \in V$ and $\mathcal{S} \subseteq T \subseteq V$. Let Υ_v be the set of nodes in V that can be reached from v with a path of length τ (possibly with self-loops): nodes in Υ_v are all and only the nodes whose load can be affected by the seed in v . We have:

$$\begin{aligned} \sigma(\mathcal{S} \cup \{v\}) - \sigma(\mathcal{S}) &= \sum_{u \in \Upsilon_v} \sqrt{\ell_{k,L}(u, \tau, \mathcal{S} \cup \{v\})} - \sqrt{\ell_{k,L}(u, \tau, \mathcal{S})} \\ &= \sum_{u \in \Upsilon_v} \sqrt{\ell_{k,L}(u, \tau, \mathcal{S}) + \ell_{k,L}(u, \tau, \{v\})} - \sqrt{\ell_{k,L}(u, \tau, \mathcal{S})}. \end{aligned}$$

For any $\beta \geq 0$, we have that $\sqrt{\alpha + \beta} - \sqrt{\alpha}$ is a non increasing function of α , and hence the last term of the previous inequality is lower bounded by

$$\sum_{u \in \Upsilon_v} \sqrt{\ell_{k,L}(u, \tau, T) + \ell_{k,L}(u, \tau, \{v\})} - \sqrt{\ell_{k,L}(u, \tau, T)} = \sigma(T \cup \{v\}) - \sigma(T).$$

We then get $\sigma(\mathcal{S} \cup \{v\}) - \sigma(\mathcal{S}) \geq \sigma(T \cup \{v\}) - \sigma(T)$ that proves the submodularity of U-BS. ◀

Then, from Theorem 3 and the above lemma, we get the following result.

► **Corollary 9.** *There exists a $(1 - 1/e)$ -approximation algorithm for the U-BS problem requiring $O(n^3(\log \tau + k))$ time.*

Proof. The result automatically follows from Theorem 3 by the monotonicity and submodularity properties of the spread function of U-BS. The greedy algorithm in Section 2.3 gives the $(1 - 1/e)$ -approximation. Computing the spread for a given seed set \mathcal{S} requires $O(n^2)$ time: loads can indeed be computed by the multiplication $A^\tau \cdot \ell_{\mathcal{S}}$, where A is the adjacency matrix of graph G and $\ell_{\mathcal{S}}$ is an n -dimensional vector encoding nodes in \mathcal{S} . Matrix A^τ is computed in $O(n^3 \log \tau)$ time with the doubling trick. Since the greedy algorithm has k iterations, and each iteration checks $O(n)$ seeds, the claim follows. ◀

5 Experiments

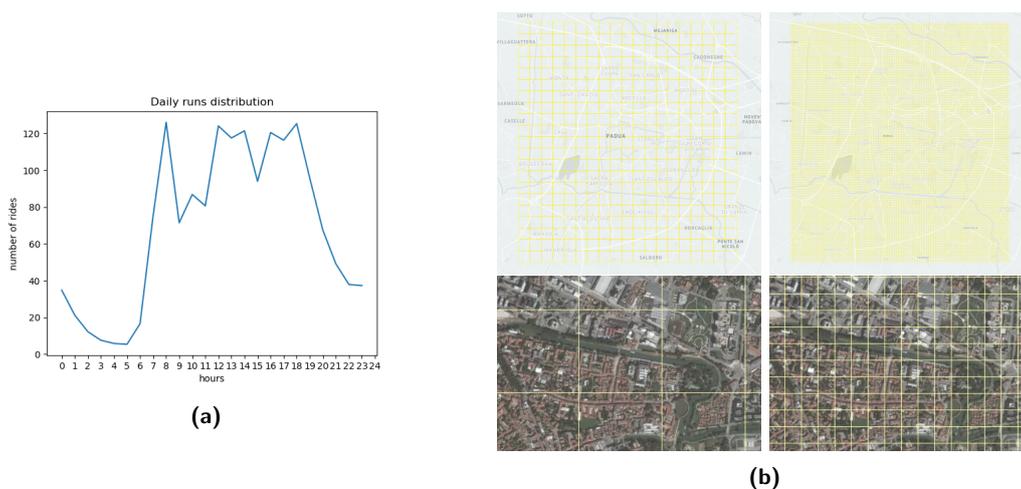
In Section 5.1, we explain how the input mobility graphs have been obtained, and then in Section 5.2 we show the findings of our analysis. The code and the input graphs are available at <https://github.com/AlgoUniPD/BikeSpreadingProblem>.

5.1 Building the mobility graphs

The input graphs used in the experiments have been built from a dataset containing all rides of the free-floating bike system in Padova (Italy) of the operator Movi by Mobike. The dataset contains 327K rides from May 1st, 2019 to January 30th, 2020. Each ride is described by the anonymized user and bike ids, and by the positions and time stamps of the pick-up and drop-off points. The graphs were constructed by following these three steps.

1. Vertexes are created by snapping each pick-up and drop-off point on a 2-dimensional grid. The grid consists of cells of size $s \times s$ with $s \in \{100\text{ m}, 500\text{ m}\}$. The two grids create two vertex sets of size 9975 and 399 respectively, covering a total area of 99.75 km^2 (see Figure 1b). Grid size should be understood as the maximum distance a user is willing to walk to find a bike.
2. The edge probabilities are constructed by two sets of rides: from 6:30 to 9:00 (morning rides) and from 16:00 to 20:00 (evening rides). In both sets, rides refer to all weekdays of October 2019. For each $u, v \in V$, the probability of an edge (u, v) is set to $n_{u,v}/n_u$, where n_u is the total number of rides originated in the cell represented by node u and $n_{u,v}$ is the total number of rides from node u to node v .
3. The graph is pruned by removing all edges with probabilities lower than a given threshold $\eta \in \{0, 0.01, 0.1\}$. For each node u , the probability mass of the removed edges originating from u is added to the self-loop (u, u) to guarantee that the weights of outgoing edges sum to one. Nodes with no edges or with only a self-loop are removed from the graph.

We thus ended up with 12 graphs $G_{s,\eta,M}$ and $G_{s,\eta,E}$ (M and E mean morning and evening, respectively), whose property is provided in Table 1. The intuition behind the three parameters (grid size s , morning or evening rides, pruning factor η) used for creating the graphs is the following. The different grid size s and pruning factor η are used for generating graphs with a different number of nodes and edges, for testing scalability. The graphs built using morning or evening rides give insight into the mobility flow during the day: the two slots are the morning and evening rush hours where workers commute to or from workplaces (see the daily rides distribution in Figure 1a). From a practical point of view, the weights computed from morning rides should be used if bikes are positioned in the seed nodes during the night to exploit the morning bike flow (equivalently, the evening rides should be used for bikes positioned in the afternoon).



■ **Figure 1** (a) Daily runs distribution during October 2019. There are three picks: around 8.00 (commuting home to work), 13.00 (lunch time, end of school), around 18 (commuting work to home). (b) Padova subdivided with a grid of 500 m and 100 m, and a detail of the subdivisions. (The underlying street and satellite maps were provided by Kepler.gl.)

■ **Table 1** Properties of the graphs used for the analysis. A graph $G_{s,r,\eta}$ with $s \in \{100\text{ m}, 500\text{ m}\}$, $r \in \{M, E\}$, $\eta = \{0, 0.01, 0.1\}$ was obtained with a grid of size s , with morning (M) or evening (E) rides, and pruning factor η .

Graph	η	n	m	avg. degree
$G_{100,0,M}$	0.0	359	1302	3.627
$G_{100,0.01,M}$	0.01	287	954	3.324
$G_{100,0.1,M}$	0.1	139	323	2.324
$G_{500,0,M}$	0.0	111	1196	10.775
$G_{500,0.01,M}$	0.01	107	1068	9.981
$G_{500,0.1,M}$	0.1	75	272	3.627

Graph	η	n	m	avg. degree
$G_{100,0,E}$	0.0	1187	5854	4.932
$G_{100,0.01,E}$	0.01	1099	4986	4.537
$G_{100,0.1,E}$	0.1	222	463	2.086
$G_{500,0,E}$	0.0	142	2625	18.486
$G_{500,0.01,E}$	0.01	125	1810	14.480
$G_{500,0.1,E}$	0.1	92	229	2.489

5.2 Performance and quality

The experimental analysis focuses on the following questions:

- (Q1) How close is the solution of the approximate algorithm to the optimal solution?
- (Q2) How does the greedy algorithm scale with input size, seed size, and step number?
- (Q3) How do the U-BS and T-BS models compare?
- (Q4) When bikes should be rebalanced?

All experiments have been executed on an Intel Xeon Processor W-2245 3.9GHz with 128GB RAM and Ubuntu 9.3.0; code was in Python 3. Running times were averaged on 3 executions.

Question Q1: exact vs approximate solutions

Table 2 shows the spread and running time for the brute force exact algorithm and the greedy approximation algorithm with the U-BS version: since the brute force has $O(n^k)$ time, which is exponential in seed size, we notice a quick increase of the running time even for small seed size. The greedy algorithm is more performing than the brute force, although

■ **Table 2** Comparison between the brute force and greedy algorithms for the U-BS version, with seed size $k \in \{2, 4\}$, $\tau = 1$, $L = 100$. The symbol * means that the instance has not be run due to excessive running time.

Graph	n/m	Brute force $k = 2$		Greedy $k = 2$		Brute force $k = 4$		Greedy $k = 4$	
		σ	time [s]	σ	time [s]	σ	time [s]	σ	time [s]
$G_{500,0.1,M}$	75/272	32.6	0.02	32.6	0.001	43.6	7.20	42.8	0.001
$G_{500,0.01,M}$	107/1068	55.3	0.07	55.3	0.002	63.9	71.36	63.9	0.005
$G_{500,0,M}$	111/1196	57.3	0.07	57.3	0.002	64.1	86.16	63.8	0.005
$G_{100,0,M}$	359/1302	121.0	1.79	121.0	0.021	173.9	18514	123.0	0.040
$G_{100,0,E}$	1187/5854	185.7	99.00	185.7	0.347	*	*	193.3	0.555

■ **Table 3** Comparison between the brute force and greedy algorithms for the T-BS version, with seed size $k \in \{2, 4\}$, $\tau = 1$, $L = 100$, $\gamma = 1$. The symbol * means that the instance has not be run due to excessive running time.

Graph	n/m	Brute force $k = 2$		Greedy $k = 2$		Brute force $k = 4$		Greedy $k = 4$	
		σ	time [s]	σ	time [s]	σ	time [s]	σ	time [s]
$G_{500,0.1,M}$	75/272	11	0.04	11	0.002	20	15.49	20	0.003
$G_{500,0.01,M}$	107/1068	31	0.22	30	0.008	35	196.93	35	0.015
$G_{500,0,M}$	111/1196	31	0.24	30	0.008	36	234.79	35	0.016
$G_{100,0,M}$	359/1302	51	8.14	51	0.089	*	*	57	0.176
$G_{100,0,E}$	1187/5854	81	302.17	81	1.03	*	*	84	2.06

■ **Table 4** Seed comparison between the brute force algorithm and the approximate algorithm for the U-BS and T-BS models for $k = 4$ (other parameters: $\tau = 1$, $L = 100$, $\gamma = 1$). The star * after a seed set means that the solution provided by the approximate algorithm has the same spread of the optimal one.

Graph	Seed overlap under U-BS	Seed overlap under T-BS
$G_{500,0.1,M}$	25%	50%*
$G_{500,0.01,M}$	100%*	75%*
$G_{500,0,M}$	25%	25%

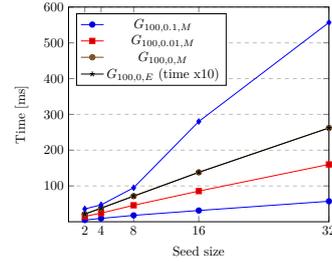
it is visible the n^3 dependence in the running time. The quality of the approximation is quite high, outperforming the theoretical worst-case upper bound of $1 - 1/e$. Similar results hold for T-BS as provided in Table 3: the running times show a small increase with respect to the previous table for U-BS, mainly due to the conditional statements for checking if a load is smaller than the threshold γ . In all cases, the spreads are very close, although the seed provided by the brute force and the greedy algorithms do not completely overlap. (see Table 4).

Question Q2: scalability

We now analyze the running time of the greedy approach for increasing values of input size, number of steps, and seed size. Since the greedy algorithms for U-BS and T-BS are almost equivalent, we only provide results for the first one. Table 5 shows the running time on the largest graph ($G_{100,0,E}$), three different seed sizes $k \in \{2, 4, 8\}$ and three different step numbers $\tau \in \{1, 10, 100\}$. For a given k , the times are almost equivalent: τ only affects the initial computation of A^τ , where A is the adjacency matrix of the graph; since the powering requires time $O(n^3 \log \tau)$, the logarithmic dependence on τ is negligible and hidden by the

■ **Table 5** Running time using $G_{100,0,E}$ with different values of the seed set size k and of the step number τ .

Graph	k	$\tau = 1$ [ms]	$\tau = 10$ [ms]	$\tau = 100$ [ms]
$G_{100,0,E}$	2	334.2	352.4	384.4
$G_{100,0,E}$	4	629.8	651.0	657.4
$G_{100,0,E}$	8	1271.8	1277.2	1278.8



■ **Figure 2** Running time with respect to seed size k . The curve of $G_{100,0,M}$ has been scaled by a factor $\times 10$ for better fitting.

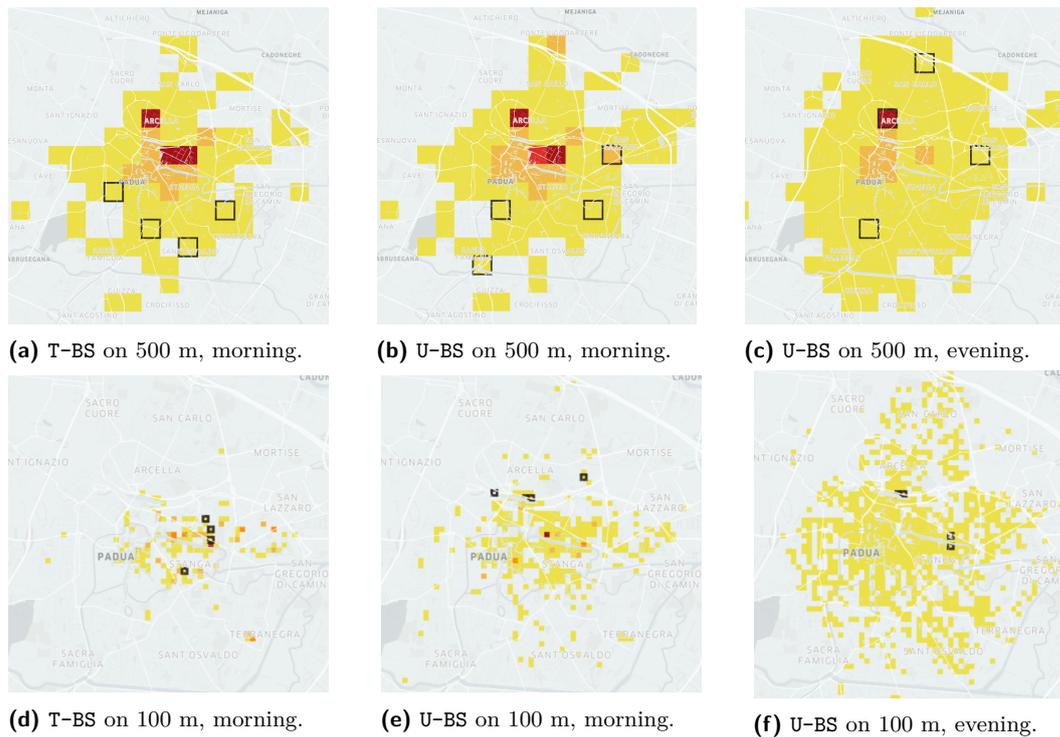
cost of the k iterations of the greedy algorithm. The previous table already shows that the running time has a linear dependency on the seed size. Figure 2 expands this analysis by considering different graph sizes. Note that the curve of $G_{100,0,E}$ has been scaled by a factor of 10 to fit the plot space. All curves show a linear dependency in k .

Question Q3: T-BS vs U-BS

In this experiment, we compare the two models. Intuitively, T-BS maximizes the number of nodes with a minimum number γ of bikes: for instance, by setting $\gamma = 1$, the algorithm maximizes the number of cells with at least one bike (in expectation). On the other hand, U-BS aims at uniformly distributing bikes among nodes, even if this implies that some nodes have a low expected number of bikes (even < 1). This allows for more fair use of bikes since it increases the load in suburb areas, differently than T-BS that facilitates central (and more crowded) areas. In the heatmaps in Figure 3, we compare the two models on the $G_{500,0,M}$ (heatmaps (a) and (b)) and $G_{100,0,M}$ (heatmaps (d) and (e)) graphs. Each heatmap shows how bike distributes, after $\tau = 2$ steps, by positioning 400 bikes in the $k = 4$ seed set selected by the greedy algorithm. We notice that U-BS covers a larger fraction of nodes in the map than T-BS. The phenomenon is more evident in the 100m grid, where U-BS colors a larger number of cells in the east and south suburb areas than T-BS. although, the majority of the cells are reached by a small number of bikes, mostly less than one bike in expectation.

Question Q4: When rebalancing?

Consider the graphs $G_{100,0,M}$ and $G_{100,0,E}$: both graphs use the 100m grid and no edge pruning. Graph $G_{100,0,E}$ has a larger number of edges and nodes than $G_{100,0,M}$ (see Table 1), highlighting a change in the use of the FFBS service from morning and evening: in the morning, rides are mostly directed towards workplaces, the train station, and university departments which are mainly located in the city center and on the east side; in the afternoon, there are much more activities (e.g., having a *spritz* with friends, going to the gym, shopping) and hence the graph covers a wider area of Padova. Bikes can be more spread in the city if we use the more vibrant afternoon for rebalancing bikes. This is confirmed by the simulation of U-BS using $G_{100,0,M}$ and $G_{100,0,E}$ in the heatmaps (e) and (f) of Figure 3: by using the evening graph, the bikes significantly spread covering also the north and west parts of the city. Similar results hold for $G_{500,0,M}$ and $G_{500,0,E}$ in heatmaps (b) and (c).



■ **Figure 3** Bike diffusion after $\tau = 2$ steps by positioning $L = 100$ bikes in each node of the seed set of size $k = 4$ selected by the greedy algorithm. For T-BS, we set $\gamma = 1$. The cells in the seed set are marked with black contours. Color scale for 500m plots: $(0.0, 2.4]$ (yellow), $(2.4, 4.8]$, $(4.8, 7.2]$, $(7.2, 9.6]$, $(9.6, 12]$ (dark red); Color scale for 100m plots: $(0.0, 1.4]$ (yellow), $(1.4, 2.8]$, $(2.8, 4.2]$, $(4.2, 5.6]$, $(5.6, 7]$ (dark red). (The underlying street maps were provided by Kepler.gl.)

6 Conclusion

In this work, we have proposed a graph approach to spread bikes in a free-floating bike system to cover a large number of zones of the service area; the idea is to select a set of zones where to position bikes and let the mobility flow spread them around the city. The current model assumes that, initially, only seed nodes have bikes while the other nodes are empty. This assumption can be removed by allowing any node to initially have some bikes: it can be shown that submodularity and monotonicity still hold for U-BS, and thus the greedy algorithm provides a $1 - 1/e$ approximate solution also in this case. Another extension to investigate is to allow a different distribution of bikes among seed nodes, to balance skewness in bike distribution due to hotspots. Finally, an important research direction is to analyze the behavior of the model on a real free-floating bike system and the differences between the theoretical findings and the actual distribution.

References

- 1 Paola Alimonti and Viggo Kann. Some APX-completeness results for cubic graphs. *Theor. Comput. Sci.*, 237(1-2):123–134, 2000.
- 2 Valentin Bouquet, François Delbot, Christophe Picouleau, and Stéphane Rovedakis. On minimum dominating sets in cubic and (claw, H)-free graphs. *CoRR*, abs/2002.12232, 2020. [arXiv:2002.12232](https://arxiv.org/abs/2002.12232).

- 3 Leonardo Caggiani, Rosalia Camporeale, Michele Ottomanelli, and Wai Yuen Szeto. A modeling framework for the dynamic management of free-floating bike-sharing systems. *Transportation Research Part C: Emerging Technologies*, 87:159–182, 2018.
- 4 Wei Chen, Alex Collins, Rachel Cummings, Te Ke, Zhenming Liu, David Rincón, Xiaorui Sun, Yajun Wang, Wei Wei, and Yifei Yuan. Influence maximization in social networks when negative opinions may emerge and propagate. In *Proc. of the 11th SIAM Int. Conf. on Data Mining (SDM)*, pages 379–390, 2011.
- 5 Iris A. Forma, Tal Raviv, and Michal Tzur. A 3-step math heuristic for the static repositioning problem in bike-sharing systems. *Transportation Research Part B: Methodological*, 71:230–247, 2015.
- 6 Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- 7 Ashish Kabra, Elena Belavina, and Karan Girotra. Bike-share systems: Accessibility and availability. *Management Science*, 66(9), 2019.
- 8 David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proc. 9th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD)*, pages 137–146, 2003.
- 9 William Ogilvy Kermack, A. G. McKendrick, and Gilbert Thomas Walker. A contribution to the mathematical theory of epidemics. *Proc. Royal Society of London, series A*, 115(772):700–721, 1927.
- 10 Yuchen Li, Ju Fan, Yanhao Wang, and Kian-Lee Tan. Influence maximization on social graphs: A survey. *IEEE Trans. Knowl. Data Eng.*, 30(10):1852–1872, 2018.
- 11 Zhi Li, Jianhui Zhang, Jiayu Gan, Pengqian Lu, and Fei Lin. Large-scale trip planning for bike-sharing systems. In *Proc. 14th IEEE Int. Conf. on Mobile Ad Hoc and Sensor Systems (MASS)*, pages 328–332, 2017.
- 12 Lei Lin, Zhengbing He, and Srinivas Peeta. Predicting station-level hourly demand in a large-scale bike-sharing network: A graph convolutional neural network approach. *Transportation Research Part C: Emerging Technologies*, 97:258–276, 2018.
- 13 Russell Meddin, Paul DeMaio, Oliver O’Brien, Renata Rabello, Chumin Yu, Rahil Gupta, and Jess Seamon. The Meddin bike-sharing world map. Accessed August 17th, 2021. URL: <http://bikesharingworldmap.com/>.
- 14 Stephen J. Mooney, Kate Hosford, Bill Howe, An Yan, Meghan Winters, Alon Bassok, and Jana A. Hirsch. Freedom from the station: Spatial equity in access to dockless bike share. *Journal of Transport Geography*, 74:91–96, 2019.
- 15 George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions – I. *Math. Program.*, 14(1):265–294, 1978.
- 16 Aritra Pal and Yu Zhang. Free-floating bike sharing: Solving real-life large-scale static rebalancing problems. *Transportation Research Part C: Emerging Technologies*, 80:92–116, 2017.
- 17 Svenja Reiss and Klaus Bogenberger. Validation of a relocation strategy for Munich’s bike sharing system. *Transportation Research Procedia*, 19:341–349, 2016.
- 18 Muhammad Usama, Yongjun Shen, and Onaira Zahoor. A free-floating bike repositioning problem with faulty bikes. *Procedia Computer Science*, 151:155–162, 2019.
- 19 Yue Wang and W.Y. Szeto. Static green repositioning in bike sharing systems with broken bikes. *Transportation Research Part D: Transport and Environment*, 65:438–457, 2018.
- 20 Yongping Zhang and Zhifu Mi. Environmental benefits of bike sharing: A big data-based analysis. *Applied Energy*, 220:296–301, 2018.