S. Mijnster

Simulation-based optimization for rebalancing the fleet of vehicles in free-floating shared mobility systems

TU Delft

felyx

Simulation-Based Optimization for Rebalancing the Fleet of Vehicles in Free-Floating Shared Mobility Systems

by

Stijn Mijnster

Master Thesis

in partial fulfilment of the requirements for the degree of

Master of Science
in Mechanical Engineering

at the Department Maritime and Transport Technology of
Faculty Mechanical, Maritime and Materials Engineering of Delft University of Technology
to be defended publicly on Friday February 17, 2023 at 10:30

Student number:	4371828	
MSc track:	Multi-Machine Engineering	
Report number:	2023.MME.8761	
Supervisor:	Bilge Atasoy	
Thesis committee:	Bilge Atasoy	TU Delft committee chair, 3mE
	Frederik Schulte	TU Delft committee member, 3mE
	Ciğdem Karademir	TU Delft committee member, 3mE
	Hein Harlaar, Andi Aliko	Company supervisors, felyx
Date:	February 5, 2023	

Confidentiality

Preface

During my master's in Multi-Machine Engineering, I developed a strong interest in logistics and data. Towards the end, I was given the opportunity to direct the scope of my thesis.

Over the past few years, I have seen shared mobility appearing more and more around me. What started with me using a shared bike because my own bike had a flat tire, evolved into using various types of shared vehicles during for example city trips or to reach destinations that are not accessible by public transport. That's when I also noticed the dark green Felyx mopeds, which allowed you to easily and quickly get from A to B or show someone around town in a fun and sustainable way. With all these shared mopeds that keep track of all the trips that are made, there must be an interesting thesis project, I thought. And so it happened, I contacted Felyx and was able to come by to discuss the possibilities. Together with some employees from the Data & Analytics team, we set up the assignment, which I presented at the TU Delft. Fortunately, it was approved, and I could start.

And now, almost a year later, I have reached the end of this whole process. And yes, I'm still smiling. Despite it being a long journey, I have learned a lot and developed myself in many aspects. Aspects such as problem-solving, analytical thinking, communication, collaboration, planning, academic writing, and so on. I also gained a thorough understanding of how things work in a fast-growing scale-up, which was one of my predetermined goals. I look back on the past year with pride, which was made possible by a few people in particular.

First of all, I would like to thank my thesis supervisor Bilge Atasoy. During our approximately bi-weekly meetings, you always made time for me and clearly explained things till I fully understood. It still amazes me, with all the students you supervise and your other responsibilities, how well you were always aware of all the details of my project. You provided constructive feedback, and every time I walked out of your office, I felt optimistic and ready to keep on going. I cannot repeat enough how grateful I am for this. I would also like to thank Frederik Schulte, my second supervisor, for your contribution and your critical view during the moments we worked together.

Then a big thank you to all the employees at Felyx who made my time here unforgettable, especially the Data & Analytics team. The daily stand-ups, team check-ins, deep dives, and of course the team events always kept me motivated. I felt like part of the team and really appreciate all the help I received.

Finally, I would like to thank my friends and family for all the support and encouragement. Listening to what I was working on and all the challenges I faced definitely helped me in the process.

I will continue to pursue my interest in logistics and data, as well as keep a close watch on the developments in the shared mobility market, as I am convinced that this can improve the quality of life in cities, which I value much.

Stijn Mijnster
Rotterdam
February 2023

Abstract

In recent years, shared mobility systems have had a growing presence in cities all over the world. This is understandable given its numerous advantages such as the reduced need for personal vehicle ownership, reduced traffic congestion and emissions, increased parking efficiency, and cost savings for users. Overall, shared mobility systems offer the potential to revolutionize transportation, providing individuals with more options and helping to create more sustainable, livable cities. For shared mobility systems to fully deliver their benefits, vehicle availability must be maintained at the right place and time. If the vehicle distribution is not optimal, it may lead to overcrowding and shortages which in turn will discourage usage and lead to reduced revenues for the operator. Therefore, ensuring proper balancing of supply and demand is crucial for the success of the shared mobility service. One way to balance supply and demand is through physically rebalancing vehicles within the service area. In this study, a simulation-based optimization model is created and used to determine the optimal rebalancing operations while quantifying system improvement. A case study is conducted using real data from the Dutch moped sharing provider Felyx to examine the impact of performing rebalancing operations in Eindhoven throughout May '22. The results demonstrate a potential increase in profit of up to 2.06%. By performing the recommended rebalancing actions several times a week in each city where the operator is active, a significant amount of extra profit can be made. This additional profit will even rise as the usage of shared mobility rises in general.

Keywords: Shared mobility · Operator-based rebalancing · Free-floating · Discrete-event simulation · Simulation-based optimization

Contents

List of figures

List of tables

1 Introduction

The shared use of vehicles, also known as shared mobility, has grown significantly in recent years. Shared mobility operators pop up all over the world, particularly in larger cities, providing a variety of vehicles to enable users to gain short-term access to transportation modes on an as-needed basis (Shaheen and Chan, **2016**). Although the proliferation of tech-enabled shared mobility systems has occurred mostly within the last decade, shared mobility services are not a new phenomenon. Examining the history of shared mobility reveals some notable events, which are highlighted in this section. Additionally, shared mobility includes a variety of terms and characteristics that are used throughout literature. A clarification of most of those terms and characteristics is given. Finally, one of the most critical challenges within shared mobility, on which this study focuses, is discussed.

1.1 Background

In addition to a few previous initiatives or trials, the first bike-sharing program in the Netherlands was launched in Amsterdam in 1965 as "the white bicycle plan" (Witte Fietsenplan), where a small number of white painted bikes were left unlocked around the city, to be used by anyone in need of transport. The concept drew a lot of attention, but it was short-lived. The free white bikes were quickly removed by the police, but this was just the beginning for bike-share schemes. It was also supposed to be a statement against the increasing number of cars, which was a major issue in Amsterdam (Davis, **2014**).

Around 1974, one of the activists (Luud Schimmelpennink) of the white bicycle plan realized one of the world's first technology-based car-sharing projects, known as the "Witkarren", which was a system for sharing small electric cars. To use a Witkar, one had to sign up as a member and pay a nominal fee every kilometer driven. The Witkarren could be driven between stations where they were charged when not in use. Due to financial shortcomings, the network of vehicles and stations was not expanded further, although there were Witkarren on the streets for ten years, proving that such a system could function (Ploeger and Oldenziel, **2020**).

Later, several initiatives entered the market worldwide, such as free bike-sharing programs with a deposit or paid options using a chipcard with pickup and drop-off stations. Eventually, the huge and unexpected success of the "Parisian bike-sharing program" in 2007 encouraged cities all over the world to establish their own systems, all modeled on Schimmelpennink's [4].

With the founding of UberCab (now known as Uber) in 2009, a new era began in terms of digitization and technology [5]. Uber started as a ride-hailing company that primarily provided on-demand car trips, but has since evolved into a Mobility as a Service (MaaS) platform that facilitates end-to-end journeys. Via this platform, for instance, users are now able to request car rides and rent bicycles, all as a single charge managed through a single user account. With this, Uber eliminates the friction of having multiple apps and providers, and is therewith dominating the mobility sector. They are even collaborating with public transport operators to include their services in the Uber app [6].

Aside from bike and car sharing, other vehicle sharing systems have entered the market. Scoot Networks, which began operations in 2011, was the first to provide electrically powered mopeds in San Francisco, which consumers could rent through a smartphone app [7]. In the Netherlands, Felyx was the first company to serve the streets with shared electric mopeds, first deploying in 2017 [8]. The company Bird was the first to launch electric scooters, which did no longer needed to be powered by foot like the original scooter [9].

With an increasing demand for shared mobility, more initiatives will undoubtedly emerge in the future. Rumors suggest that research is already underway into Urban Air Mobility (UAM), which would enable shared air transport between suburbs and within cities as early as 2025 [10].

One thing that all vehicle sharing companies have in common is that practically all vehicles have been electrically powered in recent years. Fleet managers ensure the operationality of the fleet by charging the cars at charging stations or replacing the batteries of the mopeds, bicycles, and scooters when they run out. As a result, the shared mobility provider contributes to a sustainable future [11].

1.2 Terminology

Ride-sharing, vehicle-sharing, and ride-hailing

Ride-sharing, also known as carpooling and vanpooling, occurs when passengers and drivers with similar origin-destination pairs share a vehicle. Actually, it is more accurate to say that the route is shared rather than the vehicle. During these trips, it is possible that additional stops will be made along the route to pick up and drop off passengers. Ride-sharing reduces the need for multiple cars on the road. When multiple drivers share a single vehicle over time, this is referred to as vehicle-sharing. The vehicle might be owned by a person who lends his or her vehicle to someone else (peer-to-peer), but in most shared mobility systems, an organization offers a fleet of vehicles that can be rented. The organization typically provides insurance, fuel, parking, and maintenance. Two examples are car-sharing and bike-sharing. Passengers who use ride-hailing pay a personal driver to take them to their destinations. Although this might seem similar to ride-sharing, many people who use ride-hailing services are not ride-sharing. The hired drivers don't have the intention to go to the same destination as their passengers, but instead take multiple routes to fulfill their customers' needs. These drivers almost never take on extra passengers along the route, making it a ride-hailing system rather than a ride-sharing system. A more detailed explanation of the definitions of these shared mobility terms is provided in the work by Shaheen and Cohen (**2020**). Examples of these different types of shared mobility are presented in Figure 1.

(a) Ride-sharing [13] (b) Vehicle-sharing [14] (c) Ride-hailing [15]

Figure 1: Different types of shared mobility.

Mobility on Demand and Mobility as a Service

Mobility on Demand (MoD) and Mobility as a Service (MaaS) are sometimes used interchangeably throughout literature, however, they are different. MoD focuses on the commodification of passenger mobility, goods delivery, and transportation systems management, whereas MaaS is primarily concerned with passenger mobility aggregation and subscription services. A distinguishing characteristic of MaaS, according to Shaheen and Cohen (**2020**), is brokering travel with suppliers, repackaging, and reselling it as a bundled package.

Public transport

Public transport is also a form of shared mobility, since the users of busses, trains, trams, and subways share their rides with other passengers. Public transport is known for its fixed routes and schedules.

Micromobility

Micromobility is a subgroup of mobility in general that refers to a range of small, lightweight vehicles that operate at speeds typically below 25 km/h (15 mph) and are driven by users themselves. Bicycles, e-bikes, electric scooters, and electric skateboards are examples of shared mobilitity vehicles. They usually only transport one person at a time. However, the definition has evolved and currently electrically driven mopeds with a top speed of 45 km/h (28 mph) are gradually included in this category. An overview of recent studies on several areas of micromobility is presented by Jiangping et al. (**2022**), including aspects such as type of users, travel characteristics, usage and performance, multimodal integration, competition, and methodologies to operate more efficiently.

All of the above sharing models have one thing in common; they allow customers to obtain transportation services on an as-needed basis. This is pretty typical for shared mobility systems nowadays and helps to reduce vehicle ownership. Technological developments and the advance of digitisation are the driving forces behind the emergence and rapid growth of so-called sharing platforms, which allow users to reserve the shared vehicles. However, increased urbanisation and environmental consciousness are also important influences. And lastly, according to Basselier et al. (**2018**), financial motives may likewise be contributing to this development, both among consumers and suppliers of goods and services.

1.3 Characteristics

One-way trips or round-trips

The trips of vehicle sharing systems can be classified as one-way trips or round-trips. One-way trips have their starting and ending points at different locations, whereas round-trips, possibly with a stopover, return to their original starting point of the trip. The pickup and drop-off locations for one-way trips may either be limited to a rental station or may be anywhere within an operational area in what is called a 'free floating sharing system', as described in the next paragraph. One-way vehicle sharing systems are generally more expensive to operate than traditional round-trip vehicle sharing systems because the locations from which the vehicles can be used change depending to prior users' itineraries. The optimal distribution of where the vehicles are located throughout the operational area will therefore fluctuate over time and may become unbalanced. There are several ways to return the fleet to its ideal distribution, as described in Section 1.4, but most of them are costly. Yoon et al. (**2017**) conducted a survey to investigate the factors that influence the utilization of one-way and round-trip car-sharing in Beijing. Results of the survey indicate that parameters such as gender, age, income, car ownership, and demography have significant effect on the utilization of car-sharing systems.

Station-based or free-floating

Vehicle sharing systems can be characterized as either station-based or free-floating. Station-based, as the name suggests, means that the vehicles need to be picked up at and returned to a station. In most sharing systems it doesn't matter at which station the vehicles are picked up or returned, and there can be multiple stations within a city. As an example, you are able to rent a bike in the city center and then later park it at a station near your hotel. In a free-floating vehicle sharing system, as opposed to a station-based one, the vehicles are independent of a station and can be parked anywhere within the operational area of the provider.

Vehicles in free-floating systems are equipped with GPS tracking and GSM modules, allowing operators to collect real-time positioning and riding trajectory data automatically. Examples of these different types of shared mobility systems are presented in Figure 2.

(a) Station-based [19]

(b) Free-floating [20]

Figure 2: Different types of shared mobility systems.

1.4 Matching supply and demand

From the operator's perspective, managing a shared mobility fleet involves a few interesting challenges, but there's one in particular that keeps the fleet operators occupied: maintaining a constant supply of vehicles in the right places (Trautmann and Gnägi, **2022**). Conveniently located vehicles are a defining feature of the service's user experience. Ideally, a user should always be able to find a vehicle nearby. But unfortunately for operators, the vehicles won't perfectly distribute themselves. Users often pick up vehicles at busy locations and park them in low-demand areas. This makes them out of reach for most other users, reducing their usage. If left unchecked over a period of time, the fleet as a whole can become imbalanced, with an excess of vehicles in low-demand areas and a deficiency in high-demand areas [22]. Some degree of imbalance is unavoidable. Every time a user makes a ride, they shift the spatial distribution of supply, and typically not towards the spatial distribution of demand. Operators should endeavor to minimize it, since a severe fleet imbalance can have overwhelming negative effects on ridership and revenue, eventually making it unfeasible to continue operating [23].

There are various solutions to this problem, including both hardware- and software-driven approaches. The most straightforward solution to any demand problem is to simply increase supply. Adding more vehicles to the fleet is likely to mean more will be available where and when they're needed. However simply increasing the number of available vehicles in the fleet without optimizing the distribution, may quickly turn into loss of capital and decrease of profitability (Alvarez-Valdes et al., **2016**). Another solution is by redistributing the fleet across the city, also known as 'rebalancing', where vehicles are picked up in low-demand areas and moved to high-demand areas (Pal and Zhang, **2017**). This is often done with a van or a trailer that can transport a reasonable amount of vehicles at once. Manual relocations are fairly expensive, so they should only be carried out when unavoidable or when they'll generate a net-positive improvement on the performance of the fleet to justify the cost. A better distribution of the fleet can also be achieved by incentivizing users, eliminating the need for more vehicles or manual relocations. Such as by creating pickup and drop-off incentive zones in areas with low and high demand, respectively, or by discounting certain rides to make them more appealing (Zhang et al., **2019**).

Data-driven algorithms can be used to determine the optimal number of vehicles the fleet should have, where vehicles should be picked up and moved to, or where incentives and discounts should be applied. For the past decade, researchers have investigated this area of study extensively, trying to improve these algorithms to be faster and more accurate (Mourad et al., **2019**). To contribute to this field of research by developing an algorithm that reduces the supply and demand imbalance, one must first comprehend the current state of the fleet. Additionally, in order to properly configure the algorithm's parameters, a thorough understanding of the user's behavior is required (Shaheen et al., **2017**).

1.5 Scope of this research

In this study, focus is on rebalancing the vehicles by physically relocating the vehicles throughout a city. The following research question and sub-questions are addressed:

Research question:

- *How does the implementation of simulation and optimization techniques affect the performance of rebalancing operations in shared mobility systems?*

Sub-questions:

- *What patterns in user behavior within a shared mobility system can be observed using real data?*
- *How can the patterns in user behavior be used to create a simulation model that accurately represents reality and allows for predictions of future events?*
- *How can the simulation results be used in an optimization model to determine the optimal rebalancing actions for a shared mobility system?*

The remainder of this paper is structured as follows. Section 2 reviews the related literature. The proposed methodologies for simulation as well as simulation-based optimization are presented in Section 3. Following that, Section 4 consists of a case study that is performed on real data from the Dutch moped-sharing company Felyx. In this section, first, a comprehensive data analysis is carried out, followed by the implementation of the methodologies stated in the previous section, along with the results. Finally, Section 5 concludes this research and provides recommendations for future research.

2 Related literature

Rebalancing the fleet of vehicles within a shared mobility system can be performed by either the operator or the user. In this study, the focus is on physically relocating the vehicles throughout a city, which is known as operator-based rebalancing. In this section, first of all, several characteristics of this type of rebalancing are explained. Subsequently, two distinct approaches from related literature are highlighted, namely rebalancing through mathematical optimization and rebalancing through simulation-based optimization. Finally, the research gap is adressed.

2.1 Operator-based rebalancing

The operation of rebalancing the fleet of vehicles varies depending on the shared mobility system. In ride-hailing systems, such as Uber, drivers are directed to move to high-demand areas. The drivers drive to these areas themselves. In car-sharing systems, where users themselves are the drivers, the operator picks up the vehicles and parks them in other areas. In other sharing systems, such as for mopeds, bikes, and scooters, multiple vehicles are loaded into a van or on a trailer and are dropped off at the desired locations. Rebalancing by the operator can be done several times per day at fixed intervals or whenever it is needed.

Static or dynamic rebalancing

Rebalancing by the operator can be performed in a static or dynamic system. A static system means that the vehicles are not used, therefore the locations of the vehicles do not vary during the rebalancing operations. This is most common at night when system usage is at its lowest. Practical advantages here include little or no traffic and less parking issues. In a dynamic system, the vehicles are used during the rebalancing operations, as it is throughout the day. The distribution of the vehicles changes continuously, which makes it much harder to perform the rebalancing actions since it includes a scheduling component based on the users' activity.

Reiss and Bogenberger (**2017**) evaluated both operator-based and user-based rebalancing strategies and point out that at least a share of the rebalancing tasks can be completed by users, almost cost-neutral. However, there is a certain threshold where rebalancing becomes too critical, and an operator-based intervention is unavoidable.

2.2 Rebalancing through mathematical optimization

In order to determine optimal rebalancing actions in a shared mobility system, knowledge of the demand, which varies depending on time and location, is required. Additionally, a thorough understanding of the users' behavior is crucial to gain insights into how the distribution of the demand eventually shifts towards the distribution of supply. In other words, when and where do users start their rides and where do they go to. In the literature discussed below, various strategies to approach this demand are explored, all to produce the most realistic portrayal of reality possible.

The first studies in literature on rebalancing by the operator are about static station-based bike-sharing systems because they were the first to enter the market. Later on, several studies were conducted on also dynamic station-based bike-sharing systems. Following that, both static and dynamic systems were investigated in a free-floating system as well. It is decided to concentrate mostly on literature on rebalancing within bike-sharing systems. The reason for this is that most research has been conducted on this topic, the market is therefore most developed, and there is a lot of overlap with moped and scooter-sharing systems when it comes to rebalancing.

Static station-based bike-sharing

Raviv et al. (**2013**) were probably the first to study the static case. These authors present two Mixed Integer Linear Program (MILP) models where multiple vehicles are used. In both cases, the objective is to minimize the weighted sum of the station's penalty costs and the total travel time. The penalty function for each station may represent any objective of the operator, such as the expected number of shortages. The solution is defined as the routes for each vehicle and the number of bicycles to load or unload at each station along its route. The demand in this case is based on past demand on similar days, acquired from a historical data set. The two MILP formulations are capable of solving problem instances of a moderate size of up to 60 stations with acceptable optimality gaps. Forma et al. (**2015**) later present a three-step matheuristic (combination of mathematical programming model and a metaheuristic) for the same problem. First, the stations are clustered by using a specialized saving heuristic. Following that, the rebalancing vehicles are routed through the cluster while tentative inventory decisions are made for each individual station. Finally, the original rebalancing problem is solved. The second and final step are formulated as MILP models that are solved by a commercial solver. Here, similar to Raviv et al. (**2013**), the sum of the penalties incurred at all stations and the total travel time is minimized. This method outperformed earlier methods in the literature when evaluated on instances of up to 200 stations and three reblancing vehicles.

Schuijbroek et al. (**2017**) worked on a simplified version of the models of Raviv et al. (**2013**). These authors propose a cluster-first route-second heuristic to rebalance the inventory, in which a polynomial-size clustering problem simultaneously considers the service level feasibility and approximate routing costs. The objective function contains only the cost calculated as the sum of the travel times. Benchimol et al. (**2011**) propose a simple method where a single truck rebalances bicycles in order to bring the inventory of each station to a predetermined value. Their objective is to minimize the routing cost. Chemla et al. (**2013**) revisited the model of Benchimol et al. (**2011**) and propose a relaxation of the problem yielding lower bounds. They present a branch-and-cut algorithm for solving the rebalancing problem with a single-vehicle, with results on instanced of up to 100 stations. Rainer-Harbach et al. (**2014**) propose an efficient local search algorithm and some variations of it for a generalization of the problem, considering the case of multiple trucks and with a target inventory value that is not a hard constraint, but imposed as a penalty in the objective function. Erdoğan et al. (**2014**) propose the first exact algorithm in the context where the inventory of each station must lie in a predetermined interval. They develop and implement a Benders decomposition scheme and an branch-and-cut algorithm for this problem. Instances involving up to 50 stations are solved to optimality. The problem considered by Erdoğan et al. (**2014**) assumes that the truck visits each station at most once, whereas Chemla et al. (**2013**) allow multiple visits to the same station.

Dell'Amico et al. (**2014**) study the rebalancing problem for the case where each station has a specific positive or a negative demand. The authors consider a fleet of capacitated trucks used to redistribute the bicycles throughout the network with the objective to minimize the total routing cost. They view the problem as a one-commodity pickup-and-delivery capacitated truck routing problem. The authors propose four mixed integer linear programming formulations for the problem, which they solve by branch-and-cut. Two years later Dell'Amico et al. (**2016**) improve themselves by using a metaheuristic based on Destroy and Repair to solve the problem, which lead to an increase in solving instances with 116 stations up to 500 stations. Szeto et al. (**2016**) consider a single-truck rebalancing problem in which the objective is a weighted sum of penalties for unmet customer demand and operational time on the vehicle route. The problem is solved by an enhanced version of a local search metaheuristic called Chemical Reaction Optimization (CRO). Ho and Szeto (**2014**) solve similar problems by iterated tabu search and obtained high quality solutions efficiently. Later, Ho and Szeto (**2017**) consider a single truck variant of the model of Raviv et al. (**2013**) and develop a hybrid large neighbourhood search to solve it. This algorithm is able to solve instances involving up to 518 stations and five trucks and therewith outperform the previous matheuristic of Forma et al. (**2015**). Wang and Szeto (**2021**) propose a enhanced artificial bee colony (EABC) algorithm to solve the problem with a single vehicle only. The problem aims to design the route and loading instructions for the rebalancing vehicle such that the weighted sum of the absolute deviation from the target inventory level, the penalty caused by broken bikes at stations, and the CO_2 emissions or the rebalancing vehicle are minimized.

Dynamic station-based bike-sharing

Pfrommer et al. (**2014**) consider the routes for the rebalancing process in the case when the trucks have to react in an on-line manner to the current state of the system, based on the demand for the next 30 minutes. Research in dynamic rebalancing has not yet been fully explored because of the modeling difficulties it involves. Earlier papers deal with the case when the time-dependent demand is known in advance and the truck operations are planned off-line. Contardo et al. (**2012**) present a mathematical formulation which cannot handle medium or large instances. They therefore present an alternative modeling approach that takes advantage of two decomposition schemes, Dantzig-Wolfe decomposition and Benders decomposition, to derive lower bounds and feasible solutions in short computing times. Kloimüllner et al. (**2014**) extend their previous work on the static variant of the problem by introducing an efficient way to model the dynamic case with greedy heuristics, GRASP, and variable neighbourhood search (VNS). Computational experiments are performed on instances based on real-world data, where the model for user demands is derived from historical data. Ghosh et al. (**2017**) develop a large-scale routing model that jointly considers routing costs and future expected demand. They develop two solution methodologies, one based on a natural decomposability of the model into bicycle rebalancing and truck routing, the other based on the aggregation of stations. Zhang et al. (**2017**) propose a methodology including inventory level forecasts, user arrival forecasts, bicycle rebalancing and truck routing. The authors model the problem with a multi-commodity time-space network flow model. The model is linearized and solved heuristically in a rolling-horizon mode. Chiariotti et al. (**2018**) use a discretization of time and historical data to compute an approximation of the 'survival time' of each station in the network. Rebalancing is only performed if the gain obtained by reallocating bikes exceeds the cost of moving the rebalancing truck. Some results show that some of the issues of rebalancing systems are due to an inaccurate estimation of the demand patterns. In order to be more effective, the system should not just take into account historical data, but also current trends and weather data. Shui and Szeto (**2018**) propose a rebalancing problem that simultaneously minimizes the total unmet demand and the fuel and CO_2 emission of the rebalancing vehicle. This study adopts a rolling horizon approach to break down the proposed problem into a set of stages, in which a static bike rebalancing sub-problem is solved in each stage. An EABC algorithm to optimize the route design in each stage and a route truncation heuristic to tackle the loading and unloading sub-problem are jointly used for optimization. Datner et al. (**2019**) formulate a mathematical formulation of the inventory problem with considering the interactions among stations. They use a local search algorithm that extracts information from the dynamics observed in a simulation.

Table 1 summarizes the literature on static and dynamic station-based bike-rebalancing problems according to number of rebalancing vehicles used, number of stations that can be considered, type of algorithm, methodology, and problem objectives.

Static free-floating bike-sharing

Pal and Zhang (**2017**) present a Novel Mixed Integer Linear Program for solving the problem. The proposed formulation can handle single and multiple vehicles and also allows for multiple visits to a node by the same vehicle. They present a hybrid large neighbourhood search with variable neighbourhood descent algorithm, which is both effective and efficient in solving large-scale rebalancing problems. Liu et al. (**2018**) propose an enhanced version of Chemical Reaction Optimization (CRO) to solve the problem. The computational results demonstrate that the enhanced CRO gets better solutions than the original CRO and has potential to tackle the rebalancing problem for larger, longer rebalancing duration, and more vehicle instances. The effectiveness of this heuristic compared with traditional meta-heuristics, such as variable neighbourhood search, tabu search, and genetic algorithm, is not known. Du et al. (**2020**) considers rebalancing of normal bikes and malfunctioning bikes simultaneously, in order to realize the ideal distribution. The authors present

Table 1: Summary of station-based bike-rebalancing problem literature.

Reference	Problem Type	No. of Vehicles	No. of Stations	Solution Methodology		Objective: Minimize
Raviv et al. (2013)	S	> 1	60	E	MIP	weighted sum of total travel time and penalty cost
Forma et al. (2015)	S	> 1	200	E, H	3-Step Matheuristic	weighted sum of total travel time and penalty cost
Schuijbroek et al. (2017)	S	> 1	135	E, H	Constraint Programming and MIP	tour length
Benchimol et al. (2011)	S	1	-	E	9.5-Approximation Algorithm	total travel cost
Chemla et al. (2013)	S	1	100	E, H	Branch-and-cut with Tabu Search	total travel distance
Rainer-Harbach et al. (2014)	S	≥ 1	700	H	Greedy Heuristic, GRASP, and VNS	weighted sum of the total absolute deviation from the target number of bikes, total number of loading/unloading activities, and overall travel time
Erdoğan et al. (2014)	S	1	50	E	Branch-and-cut with Benders decomposition	travel and handling costs
Dell'Amico et al. (2014)	S	1	116	E	Branch-and-cut	total routing cost
Dell'Amico et al. (2016)	S	1	564	H	Metaheuristic based on Destroy and Repair	total routing cost
Szeto et al. (2016)	S	1	300	H	Chemical Reaction Optimization	weighted sum of unmet customer demand and operational time on the vehicle route
Ho and Szeto (2014)	S	1	400	E, H	Iterated Tabu Search	total penalty cost
Ho and Szeto (2017)	S	> 1	518	H	Hybrid Large Neighbourhood Search	weighted sum of total travel time and penalty cost
Wang and Szeto (2021)	S	1	300	H	Enhanced Artificial Bee Colony algorithm	weighted sum of the absolute deviation from the target inventory level, penalty induced by broken bikes, and generated CO_2 emissions
Pfrommer et al. (2014)	D	≥ 1	-	H	Predictive model and Greedy Heuristics	operating costs for a given service level
Contardo et al. (2012)	D	> 1	-	H	Hybrid MIP approach using Dantzig-Wolfe and Benders decomposition	unmet demand
Kloimüllner et al. (2014)	D	≥ 1	-	H	Greedy Heuristic, GRASP, and VNS	weighted sum of unfulfilled demand, absolute deviation from the target fill level, total number of loading instructions, and total drive time
Ghosh et al. (2017)	D	≥ 1	300	H	MIP with Lagrangian dual decomposition	unmet demand and routing cost
Zhang et al. (2017)	D	≥ 1	200	H	Time-space network flow approach	total vehicle travel costs and user dissatisfaction
Chiariotti et al. (2018)	D	> 1	280	H	Nearest Neighbour Heuristic	unmet demand
Shui and Szeto (2018)	D	1	180	H	Enhanced Artificial Bee Colony algorithm	unmet demand and fuel and CO_2 emissions of rebalancing vehicle
Datner et al. (2019)	D	-	300	H	Simulation-based guided Local Search algorithm	a journey dissatisfaction function

Explanation of terms:

S = Static, D = Dynamic, E = Exact, H = Heuristic, MIP = Mixed Integer Programming, GRASP = Greedy Randomized Adaptive Search Procedure, VNS = Variable Neighbourhood Search

an integer linear programming model to formulate the problem, and an effective greedy-genetic heuristic is designed to solve it for large instances. Ma et al. (2021) found that the rebalancing demand in different areas has stochastic characteristics with multiple demand scenarios. They design eight stochastic simulation-based genetic algorithms to address this problem. Zhang et al. (2022) uses real data collected by a bike-sharing company in an adaptive hybrid nested large neighbourhood search and variable neighbourhood descent with several well-designed operators to solve the problem.

Dynamic free-floating bike-sharing

Caggiani et al. (2018) first propose a spatio-temporal clusterization of the system. Following that, the authors present a Nonlinear Autoregressive Neural Network model to forecast the trend of available bikes in each spatio-temporal cluster. Finally, they present a Decision Support System aimed at the maximization of user satisfaction. Warrington and Ruchti (2019) adapt the Sparse Phase and Amplitude Reconstruction (SPAR) algorithm to this problem. They cast the rebalacing problem as a two-stage stochastic program, where the operator makes rebalancing decisions in the first stage without knowledge of the realization of customer demand, which happens in the second stage. By doing this, the difficulties of the first and second stage are separated and therefore easier to solve. Luo et al. (2021) study the rebalancing of bikes among gathering points, with explicit consideration of the collection of scattered bikes under stochastic demand within a specific area. First, they formulate a Markov Decision Process (MDP) model. Following that, they design a policy function approximation (PFA) algorithm and apply the optimal computing budget allocation (OCBA) method to search for the optimal policy parameters.

Table 2 summarizes the literature on static and dynamic free-floating bike-rebalancing problems according to number of rebalancing vehicles used, number of bikes that can be considered, type of algorithm, methodology, and problem objectives.

Table 2: Summary of free-floating bike-rebalancing problem literature.

Reference	Problem Type	No. of Vehicles	No. of Bikes	Solution Methodology		Objective: Minimize
Pal and Zhang (2017)	S	≥ 1	400	H	Hybrid Large Neighbourhood Search and VND	makespan of the fleet of rebalancing vehicles
Liu et al. (2018)	S	> 1	400	H	Chemical Reaction Optimization	weighted sum of inconvenience level, unmet demand, and total operational time
Du et al. (2020)	S	≥ 1	4760	E, H	ILP and a Greedy-Genetic Heuristic	makespan of the fleet of rebalancing vehicles
Ma et al. (2021)	S	> 1	-	H	Stochastic Simulation-based Genetic Algorithm	total cost for the rebalancing vehicles
Zhang et al. (2022)	S	1	500	H	Adaptive Hybrid Nested LNS and VND	total cost of rebalancing process
Caggiani et al. (2018)	D	1	200	H	Decision Support System, Genetic Algorithm	unmet demand and lost users
Warrington and Ruchti (2019)	D	> 1	1103	H	SPAR algorithm	total cost of rebalancing process
Luo et al. (2021)	D	1	8	-	Policy Function Approximation	unmet demand

Explanation of terms:

S = Static, D = Dynamic, E = Exact, H = Heuristic, VND = Variable Neighbourhood Descent, ILP = Integer Linear Programming, LNS = Large Neighbourhood Search, SPAR = Sparse Phase and Amplitude Reconstruction

2.2.1 Summary

Most of the papers from the literature regarding rebalancing optimization of bike-sharing are devoted to static rebalancing which assumes that the demand variation can be neglected. Less are associated with dynamic rebalancing. There are also fewer papers on vehicle rebalancing in free-floating systems than in station-based systems, because free-floating bike-sharing systems entered the market later and are more complex to model. In most papers they transform free-floating systems into station-based systems by defining virtual stations (nodes) and rebalance vehicles among these virtual stations. Some other papers divide the operating area into smaller zones (segments, clusters, partitions, regions, etc.), aggregate each zone as a virtual station, and only rebalance among the virtual stations. Furthermore, the rebalancing scenarios include single-truck and multi-truck rebalancing where the stations can be visited only once or multiple times. The target function is also variable, such as minimizing the total cost, minimizing the total travel time, or it is a combination with user dissatisfaction and penalty costs. In addition, the solution methods mainly include branch and bound methods, relaxation methods, meta-heuristics and complex heuristics. Finally, some articles employ deterministic demand, whereas others employ stochastic demand. As mentioned by Neumann-Saavedra et al. (2021), employing stochastic demand requires significantly more run-time and memory capacity, which is not always possible but is preferred to obtain more accurate solutions.

In certain literature, the effectiveness of the algorithms is tested using data from real-world bike-sharing systems. The sizes of the instances they can handle varies per study. The methodologies from the studies that can handle the largest instance size are listed below.

For the static station-based problems: Rainer-Harbach et al. (2014) apply their methodology to real data from CityBike Wien, the major public bike-sharing system in Vienna, Austria. They experiment with small, medium, and large-size instances with up to 700 stations, where rebalancing is performed by a maximum of 5 vehicles. Dell'Amico et al. (2016) evaluate their strategy with data of 10 different bike-sharing systems from across the world. These real-life data sets contain data from 150 to 564 stations, rebalanced by one single vehicle. Ho and Szeto (2017) test their large neighborhood search on 518 stations and 5 vehicles. In the dynamic station-based case: Both Ghosh et al. (2017) and Datner et al. (2019) validate their methodologies on two real world data sets from Capital Bikeshare (Washington, DC) and Hubway (Boston). The data sets include 300 stations as well as 5 vehicles that perform the rebalancing.

For the static free-floating problems: Du et al. (**2020**) evaluate their methodology on two real world data sets with up to 4760 bikes from Share-A-Bull in South Florida and Divvy in Chicago. To tackle this problem, they cluster the distribution of bikes across 476 nodes and only rebalance the bikes between these virtual stations. In the dynamic free-floating case: Warrington and Ruchti (**2019**) validate their methodology using real world data from Philadelphia's public bike-sharing scheme, which includes 1103 bikes and 102 rebalancing nodes. They even increased the data set to 400 nodes to experiment with. According to this paper, if more research on free-floating systems is required, a way to map the infinite-dimensional input to a finite decision problem at acceptable computational cost must first be discovered.

2.3 Rebalancing through simulation-based optimization

As systems get more complex, it is often difficult to obtain nice form analytical models that can be used to accurately capture its behavior. At some point, the behavior of these systems may even be referred to as a 'black box'. Simulation techniques are commonly used in these situations to evaluate the system and even compare design alternatives and identify the best design among them. However, if the number of design alternatives is very large or infinite, simulation can be both expensive and time-consuming. To overcome this problem, simulation-based optimization can be used, which is a method that in this case, can determine the best design without evaluating all design alternatives. Simulation-based optimization involves the search for those specific settings of the input parameters such that an objective, which is a function of the simulation output, is maximized or minimized. Whereas simulation models are effective in imitating reality by taking into account uncertainties and randomness, optimization models can quickly and accurately reach optimal solutions; the advantages of both worlds are now combined. Several studies where simulation and optimization techniques are used in literature are highlighted below.

Barth and Todd (**1999**) developed a queuing-based discrete event simulation model that included relocations and a number of input parameters that allowed different scenarios to be evaluated. Three ways of deciding when relocations should be performed were presented: 'Static relocation' based on immediate needs in a station; 'Historical predictive relocation', which uses knowledge of expected future demand, looking 20 minutes into the future, and 'Exact predictive relocation' that can be used if perfect knowledge of future demand is available, which is impossible in the real world. The model was applied to a community in Southern California and some measures of effectiveness were calculated. The simulation model is similar to the one in this paper, but they did not develop an optimization model or ways of combining both optimization and simulation. Later, Kek et al. (**2006**, **2009**) developed an optimization model and a simulation model, but in their work only the optimization models allow for determining the relocation operations. The simulation model is just used to evaluate the performance of the systems when the relocation operations determined by the optimization model are performed. Nair and Miller-Hooks (**2011**) continued exploring optimization methods and proposed a stochastic mixed-integer programming (MIP) model to optimize vehicle relocations, which has the advantage of considering demand uncertainty. However, they did not develop a simulation model.

Cepolina and Farina (**2012**) propose a methodology, based on the Simulated Annealing (SA) algorithm to optimize the fleet distribution of a station-based car-sharing system. The reason for this is that there is no analytical expression for the cost function, so the chances are high that a local optimum is reached instead of a global optimum, and the search space is extremely large. The methodology includes a simulation model of the proposed transport system which allows one to track the second-by-second activity of each user, as well as the second-by-second activity of each vehicle. The cost function consisting of the transport management cost (i.e. the cost of vehicles) and the cost to the customer (i.e. the total customer waiting time) is minimized by explicitly simulating the arrival of the users, the departure of the vehicles from the stations and the arrival of the vehicles at the stations. Jorge et al. (**2014**) present two independent tools that can be combined: a mathematical model for optimal vehicle relocation, and a discrete-event time-driven simulation model with several real-time relocation policies integrated. Results show that relocating vehicles, using any of the methods developed, can produce significant increases in profit. Well, the developed simulation model here is only used for evaluating the rebalancing policies. Weikl and Bogenberger (**2015**) provided relocation strategies for free-floating systems for pick up and drop-off. They combine a macroscopic relocation optimization policy of moving vehicles between zones, with a rule based heuristic for station to station relocations. They make use of a historical data analysis that generates the input for the calculation of a target vehicle distribution for different target periods. If vehicle supply and demand deviate from each other, an optimization model is used to calculate profit maximizing zone to zone relocations. The relocation strategies have been tested in a real-world setting rather than in a simulation. Deng (**2015**) developed a decision support tool to assist with determining the optimal fleet configuration of a MoD system accounting for stochastic demand and the effect of conducting vehicle distribution as part of daily operations. An optimization problem is defined to find the optimal fleet configuration is terms of minimizing cost and satisfying a certain level of service. A discrete-event simulator (DES) that includes a sub-optimization model to calculate hourly rebalancing schemes is built to estimate the performance of a given configuration. Finally, an algorithm is devised that combines Particle Swarm Optimization (PSO) and Optimal Computation Budget Allocation (OCBA) techniques to efficiently search the design and decision space.

Jian et al. (**2016**) use DES to model a station-based bike-sharing system. They tackle the rebalancing problem over bikes and docks as a simulation-optimization problem. Ideally, they would apply standard simulation-optimization methods, such as stochastic gradient-search and random search, to solve the problem, but this seems computationally infeasible. Instead, they develop heuristic search procedures that use statistics from a single simulation run in order to update the allocation of bikes and docks between stations. In each iteration they generate a trial solution and evaluate it with the DES model. If the trial solution improves the objective, then they move to that solution, otherwise they stay at the last solution. They do not claim that they find local or global optima, but instead see the value of these algorithms in the improvements they make in performance relative to that of starting solutions. Marczuk et al. (**2016**) develop several optimization models for three rebalancing policies within car-sharing systems: i) no rebalancing (baseline), ii) offline rebalancing, and iii) online rebalancing. The performance of the three policies are then evaluated using the simulation program SimMobility. Zhou et al. (**2017**) propose a car-sharing optimization problem also as a simulation-optimization (SO) problem. Here, no analytical expression of the objective function is available, hence traditional (analytical) discrete optimization algorithms cannot be used. A novel metamodel is formulated, which is based on a MIP formulation. The metamodel is embedded within a general-purpose discrete SO algorithm. The combination of the problem-specific analytical MIP with a general-purpose SO algorithm enables to address high-dimensional problems and become computationally efficient. More generally, the information provided by the MIP to the SO algorithm enables it to exploit problem-specific structural information. Hence, the simulator is no longer treated as a black box.

Gómez Márquez et al. (**2021**) develop a simulation-optimization framework to determine the bike inventory for stations in a large-scale bike-sharing system. The framework helps to optimize both the bike inventory at the beginning of the day, which is the focus of static rebalancing, and the bike inventory throughout the day, which is the focus of dynamic rebalancing. They implement several simulation-optimization methods including nested partitions (NP), interactive particle algorithm (IPA), cross entropy, and discrete simultaneous perturbation stochastic approximation (DSPSA) and find that IPA provides good solutions within reasonable computing time. Jin et al. (**2022**) propose a simulation framework for evaluating different rebalancing and maintenance strategies to model the daily operations of large-scale bike-sharing systems with docking stations. The framework can be integrated with any multi-vehicle static or dynamic rebalancing optimization model. An optimization model solved by an enhanced k-means clustering method (EKM) and an Ant Colony Optimization (ACO) algorithm is provided as an example for demonstrating such integration. Although the proposed simulation framework is developed for bike-sharing systems, it can be easily modified for modeling other transportation systems with non-floating stations (e.g. electrical bikes and scooters).

Table 3 summarizes the literature on rebalancing problems where simulation and optimization techniques are used according to the methodology, the vehicles considered, and the type of shared mobility system. The contribution of this work is mentioned as well.

Table 3: Summary of simulation-based optimization rebalancing problem literature.

Reference	Simulation model	Optimization model	Methodology	Vehicle	Type
Barth and Todd (1999)	✓	✗	evaluation by simulation	CS	station-based
Kek et al. (2006, 2009)	✓	✓	evaluation by simulation	CS	station-based
Nair and Miller-Hooks (2011)	✗	✓	-	CS	station-based
Jorge et al. (2014)	✓	✓	evaluation by simulation	CS	station-based
Weikl and Bogenberger (2015)	✗	✓	-	CS	free-floating
Cepolina and Farina (2012)	✓	✓	simulation apart from optimization	CS	station-based
Deng (2015)	✓	✓	simulation and optimization integrated	CS	station-based
Jian et al. (2016)	✓	✓	simulation and optimization integrated	BS	station-based
Marczuk et al. (2016)	✓	✓	evaluation by simulation	CS	station-based
Zhou et al. (2017)	✓	✓	simulation and optimization integrated	CS	station-based
Gómez Márquez et al. (2021)	✓	✓	simulation and optimization integrated	BS	station-based
Jin et al. (2022)	✓	✓	simulation and optimization integrated	BS	station-based
This paper	✓	✓	simulation and optimization integrated	MS	free-floating

Explanation of terms: CS = car-sharing, BS = bike-sharing, MS = moped-sharing

2.3.1 Summary

The field of simulation optimization has progresses significantly in the last decade with several new algorithms, implementations, and applications. Different approaches and algorithms are used to imitate reality as precisely as possible with the simulations, and then integrated with optimization techniques to ultimately make good decisions regarding rebalancing the fleet of vehicles. There is still plenty to discover in this field of research. It is notable that almost all of the aforementioned references from the literature focus on station-based sharing systems and only a very few are applied to free-floating sharing systems. Additionally, in some cases, simulation is just used to evaluate particular rebalancing strategies, while in other cases, simulation and optimization models are really integrated. Some of the algorithms used in the literature are tested using data from real-world bike-sharing systems. The sizes of the instances they can handle varies per study. The methodologies from the studies that can handle the largest instance size are listed below.

Jian et al. (2016) apply their methodology to real data of Citi Bike in New York City. They exclude the lower demands seen on weekends and hence only use data from 14 weekdays. The data set includes 466 stations and 6074 bikes. Zhou et al. (2017) evaluate their strategy using data from Zipcar in Boston, one of the world's leading car-sharing service providers. This real-world data set contains data from 315 stations and 894 cars. Gómez Márquez et al. (2021) initially test their methodology on real data from Ecobici, a bike-sharing company in Mexico City. Later, they also test their methodology using the data of Citi Bike in New York City, similarly to Zhou et al. (2017), and achieve relatively similar results. Due to the rapid growth of Citi Bike, the data at this time includes 620 stations and 12,500 bikes. One year later, Jin et al. (2022) validate their methodology as well on real data from Citi Bike New York City. At the time, the data consists of 858 stations and 19,506 bikes.

2.4 Research gap

In this study, focus is on rebalancing the vehicles by physically relocating them throughout the city. To solve this vehicle rebalancing problem, past research in this field has mainly been devoted to developing analytical optimization models that determine the required rebalancing actions to shift the current distribution of the vehicles to a so-called target distribution. This target distribution would better meet the demand at that moment in time. However, the distribution of the vehicles should, ideally, not just meet the demand at that time, but also the demand over a longer period of time. As systems get more complex, estimating the behavior over a longer period of time can be difficult. Nice-form analytical models are hard to define and do not accurately capture the behavior of the system anymore and these systems may even be referred to as a 'black box'. Simulation techniques are commonly used in these situations because they are much better at taking into account the intricate interactions between supply and demand. They can be used to evaluate the system and even compare design alternatives and identify the best design among them. However, if the number of design alternatives is very large or infinite, simulation can be both expensive and time-consuming. To overcome this problem, a combination of simulation and optimization techniques can be used to determine the best design without evaluating all design alternatives. According to Zhou et al. (**2017**), combining simulation models with optimization techniques, also known as simulation-based optimization, is an innovative and promising area of future research. Simulation-based optimization involves the search for those specific settings of the input parameters such that an objective, which is a function of the simulation output, is maximized or minimized. Whereas simulation models are effective in imitating reality by taking into account uncertainties and randomness, optimization models can quickly and accurately reach optimal solutions; the advantages of both worlds are now combined.

In most of the mentioned studies in Section 2.3, simulation and optimization techniques are used to determine and evaluate rebalancing operations within various shared mobility systems. In some cases, simulation is just used to evaluate particular rebalancing strategies, while in other cases, simulation and optimization models are really integrated. Additionally, in some of these studies, historical data of trips is used to determine a target distribution, which is then compared to the actual distribution of the vehicles. Rebalancing actions are then suggested to reduce this imbalance. In this study, historical data of trips is also used, but not directly to set up a target distribution. First, the historical data is used as an input for a simulation model to estimate how the system is likely to behave over a given period of time. Secondly, the occurrences in the simulation are then used in an optimization model to determine the optimal rebalancing actions. The distinction is that in this case, not only is historical data examined, but it is also used to predict future events based on the current distribution of the vehicles. The simulation and optimization models are thereby integrated. In the majority of the mentioned studies where a simulation model is used, demand prediction is employed. Predicting the demand is challenging, and a lot of information is typically lacking here. This is because unmet demand, also known as latent or censored demand, is not taken into consideration. The demand prediction is primarily based on trips that took place because there were vehicles available. Therefore, in this study it is decided to only use existing data, such as the trips and idle times that actually took place. To be more precise, in most simulation models the vehicles stand idle until demand pops up nearby a vehicle to make a ride. The vehicles in the simulation model in this study stand idle for a certain idle time, sampled from the idle time distribution associated with its location before it will make a new ride. As a result, the proposed approach enables to replicate the real system reasonably well without relying on predictions of demand.

Finally, the proposed simulation-based optimization model in this study is tested on real data and is applicable to various free-floating shared mobility systems as well as to any city where the operator have been operational for a sufficient length of time that enough data is collected, with one year being preferred.

3 Proposed methodology

This section presents the proposed methodology. First, a description of the problem is provided, along with a conceptual description of the model's intended usage. Secondly, the model is formulated, where the concept of simulation is first described, followed by a description of the concept of simulation-based optimization.

3.1 Problem description

We consider a vehicle rebalancing problem in a free-floating shared mobility system from the perspective of the operator. The following notation is used.

We assume a shared mobility operator providing shared *vehicles*; vehicles that have been modified in such a way that they can be accessed and used by anyone via an app. Examples of these vehicles are bicycles, cargo bikes, mopeds, and cars. Individuals who use these vehicles are referred to as *users*. The vehicles are provided in cities where the shared mobility operator is active. These cities have a *service area*; a GPS-based virtually confined area where the vehicles can be used. The user can make a *ride*; a ride can start and end at any location within the service area. The number of active vehicles and their related locations, we assume as the *supply*. The number of individuals who want to take a ride and their locations are referred to as the *demand*. To match the locations of the vehicles with the locations of the potential users, in other words, to match supply and demand, *rebalancing vehicles* are used. These rebalancing vehicles pick up the shared vehicles and relocate them within the service area. Vans or trailers that can transport a reasonable amount of vehicles at once are commonly used as rebalancing vehicles to perform these *rebalancing actions*.

After a user completes their ride, it leaves the vehicle at its destination. The vehicle will then be available to other users to make a ride. The elapsed time between the end of the ride and the start of a new ride for the same vehicle is referred to as the *idle time*. The idle time is linked to the end location of the first ride. The pickup location of the vehicle by the user depends partially on where the user is situated, but of course also on the supply of vehicles at that time of the day. The user's intent determines the drop-off location after a ride. The movement of the vehicle from the pickup location to the drop-off location is assumed as the *transition*. This transition has a certain duration, which we refer to as the *travel time*. This travel time depends on the route taken by the user and may vary due to traffic at that time of the day.

An example of a shared mobility system where rebalancing actions are performed is illustrated in Figure 3. The black border line denotes the service area. The grey circles depict users that want to make a ride. The green squares represent vehicles that are making a ride within the service area. These rides are indicated by a thin dotted line with a green location symbol as the start point and a red location symbol as the end point. The orange squares represent vehicles that are currently not in use and are therefore standing idle. The blue rectangle represents a rebalancing vehicle that is performing a rebalancing action, which is indicated with a thick dotted line with again a green location symbol as the start point and a red location symbol as the end point.

The goal is to reduce the imbalance between vehicle supply and demand by using rebalancing vehicles that perform rebalancing actions.

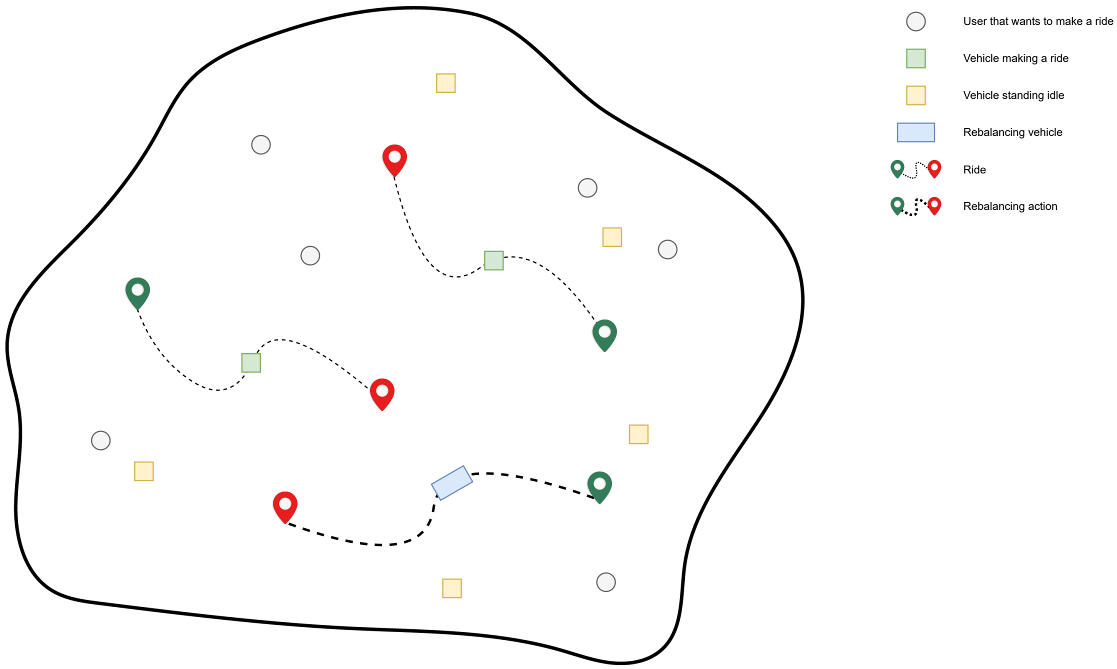

Figure 3: An illustration of a shared mobility system where rebalancing actions are performed.

Conceptual description of the model

The model's user specifies a number of parameters, including the city under consideration, the start moment of the simulation, and the desired simulation period. The start moment of the simulation can be in the past to simulate a period that has already passed and compare it to what actually happened in that period, but it can also be in the present, for example, at the start of a day. The simulator then retrieves historical data up until the start moment of the simulation. This data includes the locations of the vehicles at the start moment of the simulation and is also used to base the simulator's decisions on. After the simulation period, the simulation events are used to determine the optimal rebalancing actions using an optimization model. Finally, a comparison is made between the system with the initial distribution of the vehicles and the system with the distribution of the vehicles after the rebalancing actions. The model presents the potential improvement for various rebalancing capacities, by showing the total trips, total travel times, and extra generated revenue. The costs are then included in order to present the potential final profit.

The model may not only be used when needed but can also run automatically in the background. At various intervals, the model takes the present time as the start moment of the simulation and calculates the system's potential improvement over a certain period. If the improvement exceeds a certain threshold, a notification is sent to the user. The user is then alerted that it is now time to perform rebalancing actions in order to enhance the system and can act accordingly.

3.2 Model formulation

Below, the simulation approach is described, followed by a detailed explanation of the simulation-based optimization model.

3.2.1 Simulation

A discrete-event simulation (DES) model is created to simulate the behavior of vehicles in a shared mobility system over time. The simulation model is written with Python 3.18.13 along with the Salabim 22.0.7 Python package.

In short, the simulation model is constructed as follows. There is one class, which is the vehicle, that follows a certain process. This process consists of an initial step followed by three sequential events that are repeated till the simulation runtime is over. During the simulation, all information about these events is monitored, including the timestamps, locations, and vehicles involved. After the simulation, this information is evaluated and several key performance indicators (KPIs) are established.

The simulation model takes as input disaggregate historical reservation data from actual rides. Which data to use, is defined by the simulation settings. Furthermore, how this data is used as model inputs, how the simulation process looks like, and what the simulation outputs are, is explained in further detail below.

Simulation settings

Before starting the simulation, the simulation settings are entered. These settings include:

- The city that will be simulated
- The start date of the simulation
- The time window during the day that will be simulated
- The number of days the simulation will run
- The number of replications the simulation will run per day

Simulation model inputs

From the disaggregated historical reservation data, the model inputs are determined, as described in more detail below.

<u>Idle times</u>
The elapsed time between the end of a ride and the start of a new ride for the same vehicle is referred to as the idle time. This idle time is linked to the end time and location of the first ride. The data set for the idle times comprises data from four weeks prior to the simulation's start date and from four weeks after the simulation's start date, but from a year ago. It is expected that this data set accurately represents the idle times, with the assumption that the idle times fluctuate significantly throughout the year and therefore only these eight weeks are considered. This approach includes the trend of the previous four weeks together with the characteristics of the upcoming four weeks.

Transitions

Rides with the same start location may have different end locations. To determine the location a vehicle is likely to travel to, the various end locations are mapped. The data set for these travels, also known as the transitions, comprises data from one year prior to the simulation's start date. It is expected that this data set accurately represents the transitions, with the assumption that the transitions are generally consistent throughout the year, only depending on day of the week and time of the day.

Travel times

Rides last for a particular amount of time, which primarily relies on the start and end location of the ride. Some variation in the duration of a ride with the same start and end location may be due to a different route being taken or to traffic congestion. The data set for the travel times comprises data from one year prior to the simulation's start date. It is expected that this data set accurately represents the travel times, with the assumption that the travel times do not vary significantly during the year, during the week, or during the day.

As described above, historical reservation data is used to set up the model inputs for the simulation. This reservation data usually includes the exact location of the start and end points of rides, which are retrieved by GPS sensors located on the vehicles. The sensors' output is generally in geographic coordinates (latitude/longitude). Analyzing this data based on its exact location is both difficult and expensive. As a result, it is common practice to enclose this data in grid cells. These grid cells aggregate the underlying data points, which are then represented by a 'small' area. In this way, analyses can be carried out much easier and more efficient. There are several grid cells that can be used, each having their own application. Shared mobility systems rely on accurate mapping of geographical areas for their services. Therefore, it is crucial to use a grid map that minimizes distortion and quantization error introduced when users move through a city, which is the case with hexagonal grid cells. Uber also analyzes spatial data using hexagonal grid cells and has open-sourced it's hexagon mapping library H3, which can be used for this [72].

Simulation process

As previously mentioned, the simulation model has one class, the vehicle, that follows a certain process. This process is described in more detail here.

The initial step is to retrieve the current location of all vehicles at the start moment of the simulation. Following that, at time step zero, all vehicles are created in the simulation model and given their corresponding location and vehicle ID. As a result, the vehicle distribution is exactly the same as in reality at that point in time.

The second step in the process is that all vehicles are given a certain idle time. This idle time is location- and time-dependent and is drawn from a distribution. This distribution is based on data associated with the vehicle's location and the current time step in the simulation. As a result, the idle time for each vehicle will be unique. All vehicles will wait until the idle time is over before proceeding to the next event.

Because each vehicle has been assigned a unique idle time, the next event for each vehicle will occur at different time steps in the simulation. For now, we will focus on a single vehicle to describe the next steps in the process of the simulation.

The third step in the process, once the idle time is over, is for the vehicle to make a ride. The start location of the ride, which is the current location of the vehicle, is of course known, but the end location must be determined. The simulation model derives the end location based on data from rides with the same start location as where the vehicle is currently located and around the current time step in the simulation. It selects the end location relying on a probability distribution.

Now that the end location has been determined, the vehicle will travel from its current location to the end location of the ride. Of course, this movement takes time, which must be accounted for in the simulation. Therefore, the fourth step of the process is to establish the travel time belonging to this ride. The travel time is drawn from a distribution based on data from rides with the same start and end location. In contrast to the idle time and the end location of the ride, it is assumed here that the travel time is independent of time.

After this travel time, the vehicle will be at its new location, and steps two to four will be repeated until the simulation runtime is over. Although only one vehicle is considered here, all vehicles follow the same process simultaneously.

Simulation model outputs

During the simulation, various data is collected that contains information about the events that have occurred. First, for the idle time, the vehicle involved is tracked, as well as its location and the idle time assigned to it. Second, for the transition, the start locations and where these vehicles will travel to are monitored. Third, for the travel times, the duration of the rides is logged. At all events, the date, day of the week, and timestamp during the day are stored. Lastly, the data includes the simulation replication to which it belongs.

A set of KPIs have been defined to assess the simulation on several levels. The values of the KPIs for each simulation replication are used to determine the simulation's means and standard deviations. It is desirable to run the simulation for multiple replications to obtain a better estimate of the mean performance. The number of replications required depends on the desired accuracy and the amount of time available. The simulation's means and standard deviations can then be compared against real data. The KPIs are evaluated per day and consist of:

- The total number of rides that have taken place
- The average idle time after each ride
- The total idle time over all rides
- The total travel time over all rides

3.2.2 Simulation-based optimization

As described above, based on the location of the vehicles at the start of the simulation, the shared mobility system can be simulated over a certain period of time. To better match supply and demand, rebalancing actions can be performed. These actions alter the start location of the vehicles that are rebalanced. In other words, the distribution of the fleet of vehicles throughout the service area is adjusted. In most cases, randomly moving vehicles does not improve the KPIs of the system. Therefore, an optimization model is used to determine the optimal rebalancing actions to improve the performance of the system and thus better match supply and demand.

As input for the optimization model, data from the simulation is used. This data includes all rides taken during each simulation replication, as well as all idle times at each location after these simulated rides. The information on all rides is then used to compute the average number of outgoing and incoming rides per location per hour. With this, together with knowing the number of vehicles at each location at the start of the simulation, the evolution of the number of vehicles per location per hour can be determined. It is then possible to identify which locations have the greatest surplus of vehicles during the simulation period. Please take note that the number of outgoing rides per location per hour is limited by the number of vehicles available at this location at that moment in time. As a result, the identified surplus will never be negative for any location. To further explain this, consider the following example.

Illustrative example

At the start of the simulation, there are ten vehicles at location A. During the first hour, an average of seven rides leave this location and three rides enter. As a result, the number of vehicles after one hour is $10 - 7 + 3 = 6$. In the second hour, an average of two rides leave this location and five enter. Therefore, after two hours, the number of vehicles is $6 - 2 + 5 = 9$. It can be said that during these two hours, there was a surplus of six vehicles. Vehicles that were essentially not needed at this location. At location B, there were three vehicles at the start of the simulation. During the first hour, an average of three rides leave this location and none enter. As a result, the number of vehicles after one hour is $3 - 3 + 0 = 0$. In the second hour, no rides leave this location and only one enters. Therefore, after two hours, the number of vehicles is $0 - 0 + 1 = 1$. During these two hours there was no surplus of vehicles. It can be said that there may have been more rides that wanted to leave this location, but as there were none available, this was impossible. Therefore, it may be stated that there was a potential deficit of vehicles at this location.

For the optimization model, two scenarios have been considered. In *scenario 1*, only the idle times of the vehicles per location are taken into account, and the vehicles are rebalanced from locations with high idle times to locations with low idle times. In *scenario 2*, the idle times of the vehicles per location are also taken into account, as well as the magnitude of the surplus of vehicles at this location during the simulation. Vehicles are rebalanced from locations with high idle times and a large surplus of vehicles to locations with low idle times and a potential deficit of vehicles.

The vehicle rebalancing problem for both scenarios is defined on a directed graph $G = (H, A)$, where the set H contains the hexagons and the set A contains the arcs. The set of arcs $A = H \times H$ consists of all feasible arcs as $A = \{(i, j) \mid i \in H, j \in H, i \neq j\}$. $H1 \subset H$ and consists of all hexagons with one or more vehicles at the start of the simulation. $H2 \subset H$ and consists of all hexagons with an average idle time (determined from all simulation replications) that is based on more than 100 data points. $H3 \subset H$ and consists of all hexagons with a surplus of vehicles greater than zero. $H4 \subset H2$ and contains the hexagons with the 25% lowest average idle times. Table 4 summarises the sets, parameters, and decision variables used in both binary integer programming (BIP) formulations.

Table 4: The sets, parameters, and decision variables for the Vehicle Rebalancing Problem.

Sets	
H	Hexagons, indexed by i, j and $\|H\| = h$
A	Arcs
$H1$	Hexagons with one or more vehicles at the start of the simulation
$H2$	Hexagons with an average idle time based on more than 100 data points
$H3$	Hexagons with a surplus of vehicles greater than zero
$H4$	Hexagons with the 25% lowest average idle times and based on more than 100 data points
Parameters	
$idle\_time$	Average idle time of vehicles in a hexagon
$surplus$	Surplus of vehicles in a hexagon.
rc	Rebalancing capacity
Decision variables	
x_{ij}	Binary variable

The following decision variable is used:

The binary variable

$$x_{ij} = \begin{cases} 1, & \text{if a vehicle is rebalanced from hexagon } i \text{ to hexagon } j \\ 0, & \text{otherwise} \end{cases} \qquad \forall\,(i,j) \in A$$

The model for scenario 1 is formulated as follows:

Minimize:

$$\sum_{i \in H1} \sum_{j \in H2} \frac{idle\_time_j}{idle\_time_i} * x_{ij} \tag{1}$$

Subject to:

$$\sum_{i \in H1} \sum_{j \in H2, j \neq i} x_{ij} = 1 \tag{2}$$

$$x_{ij} \in \{0,1\} \quad \forall\, i,j \in A \tag{3}$$

The objective function in scenario 1 (1) minimizes the idle time of the pickup location divided by the idle time of the drop-off location. The decision variable then represents the optimal rebalancing action. Constraint (2) ensures that the decision variable only contains one rebalancing action. The set $H1$ ensures that vehicles can only be picked up at hexagons with vehicles available. The set $H2$ ensures that vehicles are dropped off at hexagons with an average idle time based on a sufficient number of data points. Constraint (3) defines x as a binary variable.

The model for scenario 2 is formulated as follows:

Minimize:

$$\sum_{i \in H3} \sum_{j \in H4} \frac{idle\_time_j * surplus_j}{idle\_time_i * surplus_i} * x_{ij} \tag{4}$$

Subject to:

$$\sum_{i \in H3} \sum_{j \in H4, j \neq i} x_{ij} = 1 \tag{5}$$

$$x_{ij} \in \{0,1\} \quad \forall\, i,j \in A \tag{6}$$

The objective function in scenario 2 (4) minimizes the idle time multiplied by the magnitude of the surplus of the pickup location, divided by the idle time multiplied by the magnitude of the surplus of the drop-off location. Again, the decision variable represents the optimal rebalancing action. Constraint (5) ensures that the decision variable only contains one rebalancing action. The set $H3$ ensures that vehicles can only be picked up at hexagons with vehicles available. The set $H4$ ensures that vehicles are dropped off at hexagons with the 25% lowest idle times and with an average idle time based on a sufficient number of data points. One thing to note is that if the surplus of vehicles for a particular hexagon is zero, this value is adjusted to 0.0001. This ensures that the objective function will never be zero and that therefore always a solution can be found. Constraint (6) defines, again, x as a binary variable.

In both scenarios, the optimal rebalancing action is determined. This rebalancing action causes the number of vehicles in a certain hexagon to decrease by one and to increase by one in another hexagon. The initial distribution of the vehicles over the hexagons therefore changes. This alters the calculation for determining the surplus of vehicles in the hexagons as well. By applying the change in the vehicle distribution and recalculating the surplus of vehicles in the hexagons, it is possible to run the optimization model again and determine the next optimal rebalancing action. This process can be repeated until a certain rebalancing capacity is reached. The idle time per hexagon is not dependent on the distribution of the vehicles and therefore does not change during this recalculation of the surplus of vehicles. Algorithm 1 describes this approach using pseudocode.

Algorithm 1 Determine multiple rebalancing actions

$fleet\_distribution$ = **initial fleet distribution**
$idle\_times$ = **average idle times**
$i = 0$

while $rc < i$ **do**
    $surpluses$ = **calculate surpluses**($fleet\_distribution$)
    $rebalancing\_action$ = **optimize**($surpluses, idle\_times$)
    $fleet\_distribution$ = **perform rebalancing action**($rebalancing\_action, fleet\_distribution$)
    $i += 1$
end while

4 Case study

Shared mobility systems are nowadays widely used in the major cities of the Netherlands. There are several providers offering a wide range of sharing vehicle options, including bicycles, cargo bikes, mopeds, and cars. The majority of these vehicles are electrically powered, which makes them more appealing to users while also being environmentally friendly.

This case study focuses on the shared moped provider Felyx, which is one of the four moped-sharing operators in the Netherlands at the time of writing [73]. Felyx was the first on the market when it launched in the Netherlands in 2017. Since then, Felyx has grown rapidly and now operates a fleet of over 7500 mopeds in and around 17 cities in the Netherlands, Belgium, and Germany. Every day, there are thousands of customers using the shared mopeds of Felyx.

This chapter analyzes the data of Felyx and examines one particular city in further detail. Subsequently, the methodology described in Section 3 is applied to this data and the results are presented.

4.1 Felyx data

The number of rides, as well as the start and end location of rides, are both dependent on a variety of factors. Consider for instance, people who use a Felyx moped to commute throughout the week, while others use it around town in their spare time on weekends. Aside from the user's different purposes, a moped with sufficient battery life must, of course, be available nearby, as must the weather be acceptable for riding a moped. It may be argued that this user behavior is unpredictable, yet there are patterns to be recognized. In order to notice these patterns and act on them in the future, insight into the data is required. Fortunately, Felyx has been around for some years and collected a lot of data that can be used in a data analysis.

One of the major factors that clearly influences the utilization of Felyx mopeds is seasonality. Although the weather may have an effect, individuals also tend to spend more time indoors during the winter and outside during the summer. Figure 4 shows the average number of daily rides in the Netherlands throughout the year, with a noticeable low and high season. Additionally, it is evident that the number of daily rides increased year over year. It's important to note here, that this rise is partly a result of Felyx expanding its total fleet size and the number of cities in which it operates over time.

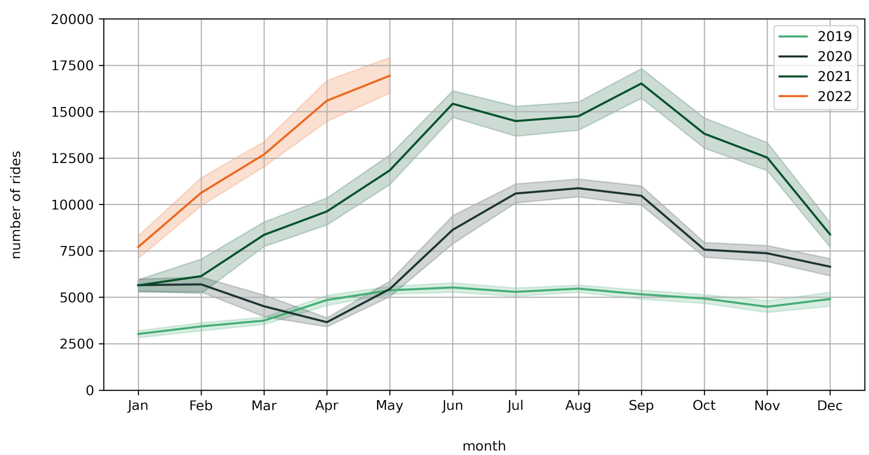

Figure 4: Average number of daily rides in the Netherlands with a 95% confidence interval.

As mentioned, weather conditions will undoubtedly play a part here as well. The Felyx mopeds are open vehicles, which means that users get wet when it rains. Also, the road surface may become slippery, causing customers to choose for a covered vehicle that is more stable on the road. Figure 5 depicts the average daily temperature in the Netherlands at 2 p.m. throughout the year, which follows a similar pattern as the average number of daily rides. Worth mentioning here is that a high temperature does not always imply the best conditions for riding a moped. Even though the temperature is high, it may rain all day or be very windy. On the contrary, on a somewhat cooler day, it might be completely dry and windless.

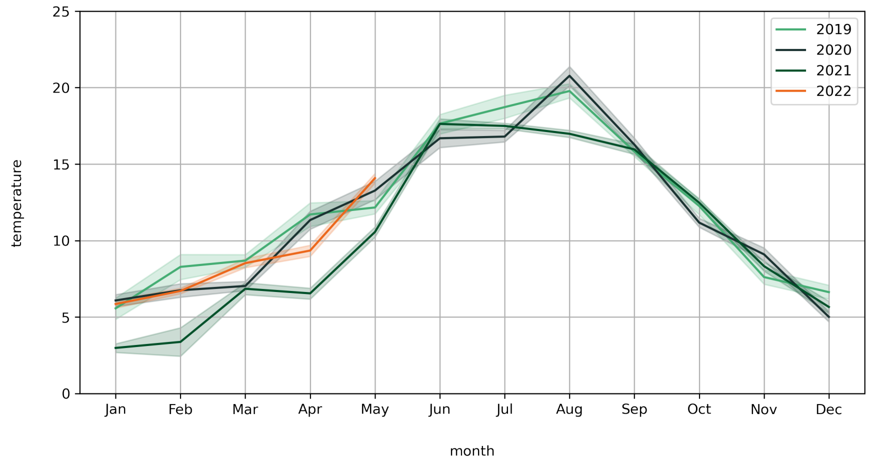

Figure 5: Average daily temperature in the Netherlands with a 95% confidence interval.

Aside from seasonality and weather conditions having an influence on the utilization of the mopeds, also other factors are at play. A comprehensive data analysis is conducted to acquire a better understanding of how the mopeds are used. Since each city has its own (spatial) characteristics, it is decided to focus on a single city. The following criteria are taken into consideration deciding on the most suitable city: 1) Felyx has been operating in this city for some time, meaning that there is data dating back at least one year; 2) enough data is being collected from a reasonable fleet size with a sufficient number of daily rides; 3) the service area of this city was altered as little as possible last year; and 4) there are as few rides coming in from or going out to surrounding cities as possible. The most suitable option out of all the cities where Felyx operates is Eindhoven.

4.1.1 Data analysis on Eindhoven

Felyx has been operational in Eindhoven since September 2020. In this data analysis, all data from this city in the year 2021 is examined.

On January 1, the fleet of Eindhoven contained 200 mopeds. In early March, another 50 mopeds were added to the fleet, bringing the total to 250 mopeds. Following that, no changes were made to the fleet size till the end of 2021.

During 2021, the service area in Eindhoven, where users could start and end their rides, looked primarily like the representation in Figure 6. Only minor changes were made to the service area if they were truly essential. This figure further illustrates that, in addition to Eindhoven, the service area also includes Veldhoven. This enabled users to travel between these two cities. The red areas indicate places where the user is not allowed to end their ride.

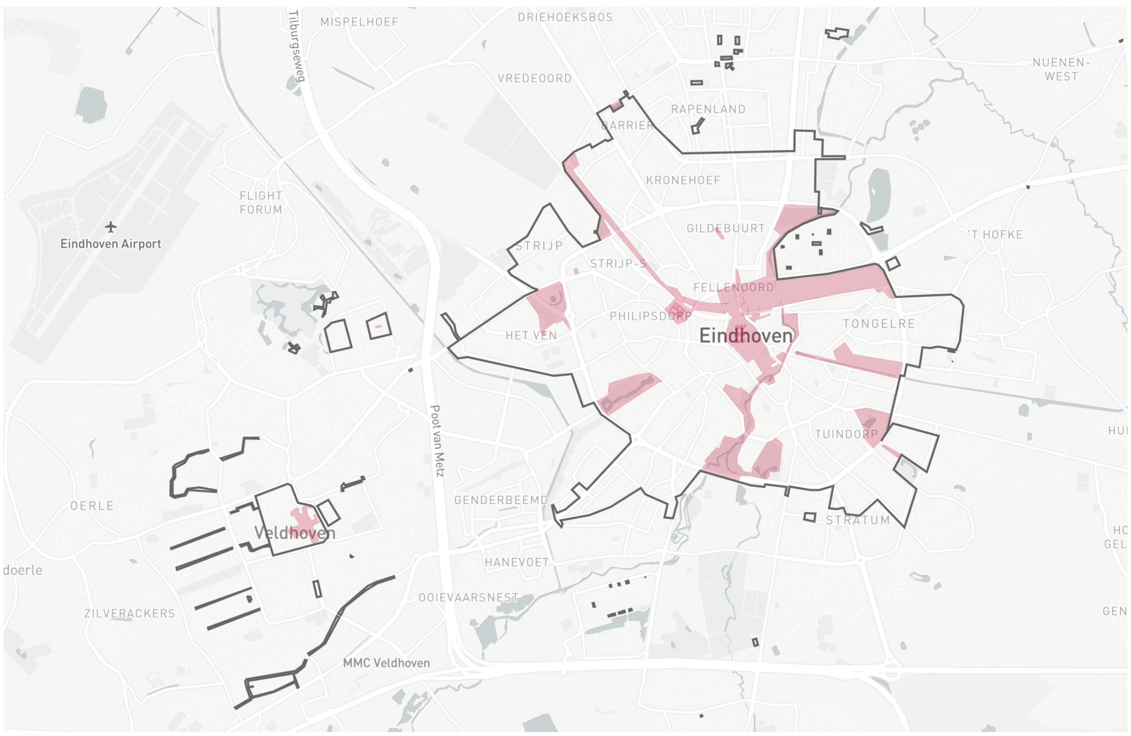

Figure 6: Representation of the service area of Eindhoven in 2021.

Instead of looking at the number of rides in all cities where Felyx operates within the Nederlands, now only Eindhoven is considered. Figure 7 shows the average number of daily rides throughout the year, again with a substantially higher number of rides in the summer months compared to the winter months. Furthermore, the number of rides at the end of the year in December is considerably higher than at the beginning of the year, implying that the utilization of Felyx mopeds has increased among users in Eindhoven.

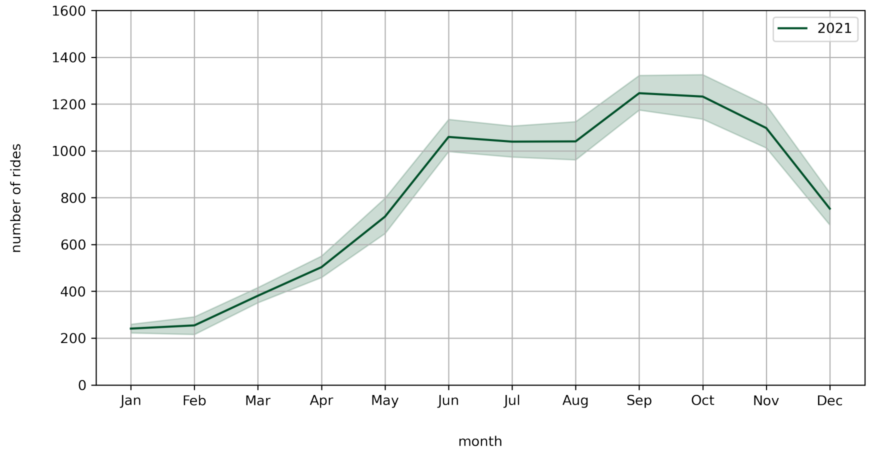

Figure 7: Average number of daily rides in Eindhoven with a 95% confidence interval.

If the number of rides is presented by the day of the week, rather than the month of the year, an interesting insight emerges, shown in Figure 8. Here, it is evident that the average number of daily rides increases as the weekend approaches, reaching a high on Saturday.

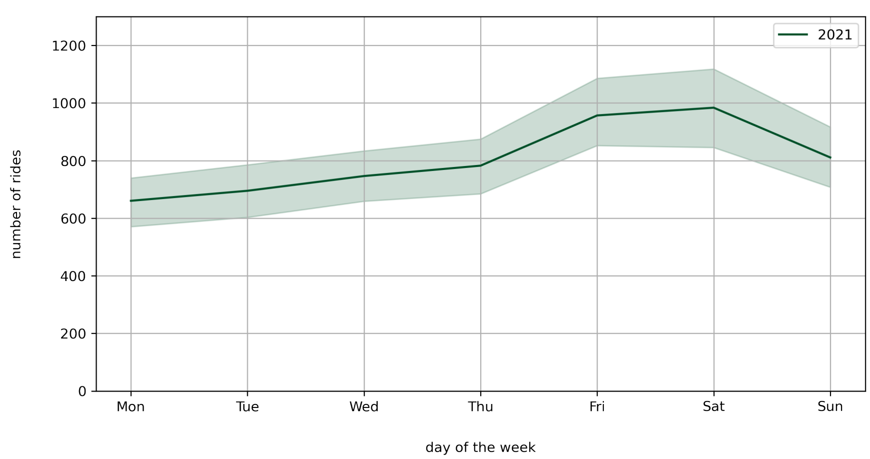

Figure 8: Average number of daily rides in Eindhoven with a 95% confidence interval.

As previously mentioned, the Felyx mopeds are used for a variety of reasons. Some customers use it for their weekly commutes, while others use it in their spare time on weekends to get around town. Of course, there are many more purposes for which the mopeds are used. In addition, some customers use the mopeds several times per week, while others only use them once a month. All these different customer behaviors make it challenging to accurately determine the demand, although patterns can be seen in the data.

Figure 9 depicts the average number of rides per hour of the day, segmented by the day of the week. What stands out here is that: 1) there are about three peaks throughout the weekdays between 08:00 and 09:00, 12:00 and 13:00, and 17:00 and 18:00; 2) weekend days differ from this weekday pattern by following more of a single wave motion that peaks around 15:00; and 3) on any day of the week, more rides take place later in the day.

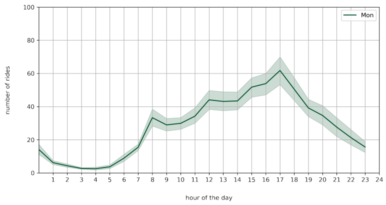

(a) Monday

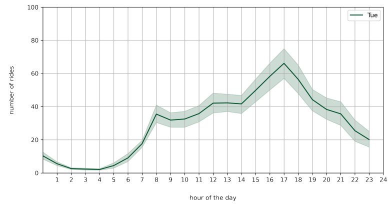

(b) Tuesday

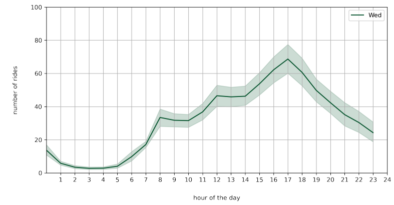

(c) Wednesday

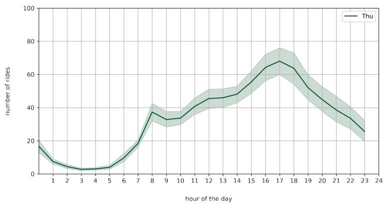

(d) Thursday

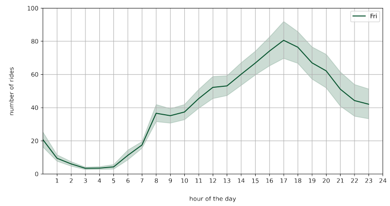

(e) Friday

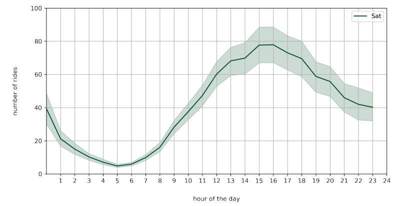

(f) Saturday

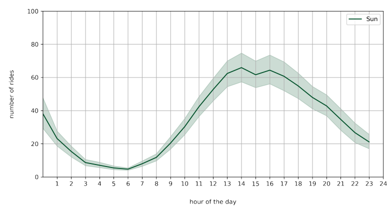

(g) Sunday

Figure 9: Average number of hourly rides in Eindhoven with a 95% confidence interval.

Certain areas of the city are more frequently used as the start and/or end points of a ride. The start location of all rides over the whole year is visualized as a heat map in Figure 10 using Kepler, an open-source application for visual exploration of large-scale geolocation data sets, created by Uber [74]. The majority of these locations are situated inside the service area (Figure 6). This makes sense, because users cannot end their rides outside of the service area, so usually rides do not start there either.

Figure 10: Heat map of the start location of all rides in Eindhoven in 2021.

As previously indicated in Figure 9, three peak moments were seen in the data during weekdays. The presence of the peak moments does not imply that the same kind of rides are made. Where rides start and end is dependent on the demand and supply for that specific moment, which is often referred to as ride patterns having spatial and temporal dependencies. Figure 11 depicts heat maps of the start location of all rides during the peak moments on a randomly chosen day. Figure 12 shows the same data, but then as rides with their corresponding start (green dot) and end (red dot) locations. These two figures illustrate that there are indeed variations in the rides during the peak moments. Further investigation reveals that these variations exist not only during the peak moments, but throughout the whole day.

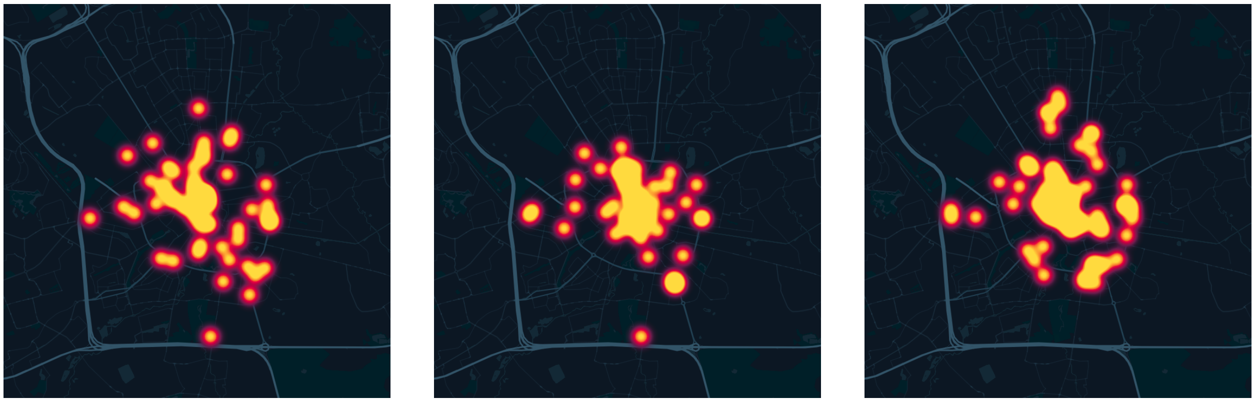

(a) 08:00 - 09:00 (b) 12:00 - 13:00 (c) 17:00 - 18:00

Figure 11: Heat maps of the start location of all rides in Eindhoven during the peak moments on July 1st.

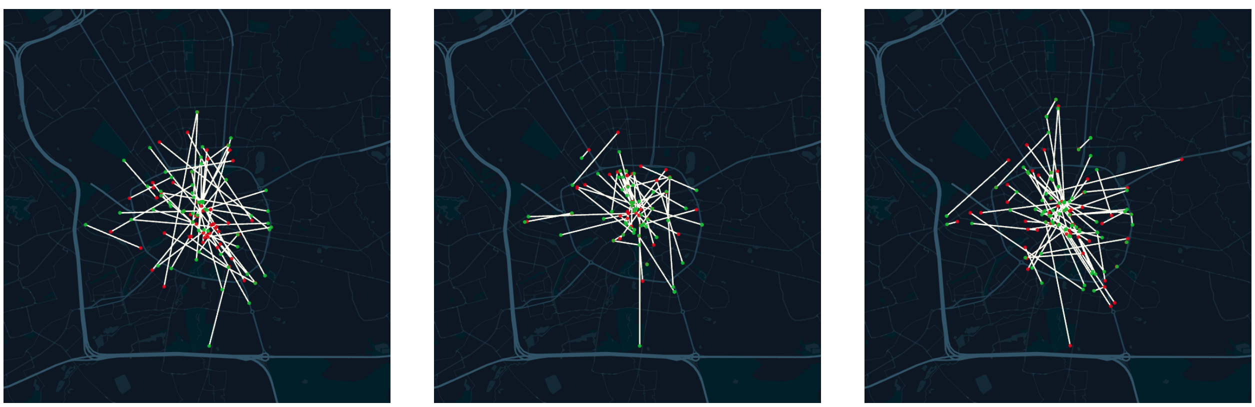

(a) 08:00 - 09:00 (b) 12:00 - 13:00 (c) 17:00 - 18:00

Figure 12: Visualization of all rides in Eindhoven in 2021 during the peak moments on July 1st.

Users travel various distances if they have different start and end points for their rides. Furthermore, even if a ride has the same start and end points, not every user takes the same route. This variation is also seen in the duration of the rides, which, in addition to taking different routes, may also be affected by factors such as traffic at that time of the day. Figure 13a depicts the distribution of the daily traveled distance per ride, with the peak implying that most rides are 2 kilometers in length. The distribution of the daily duration per ride is shown in Figure 13b, with the peak indicating that the majority of the rides last 5 minutes. The kilometers and minutes are rounded to the nearest integer.

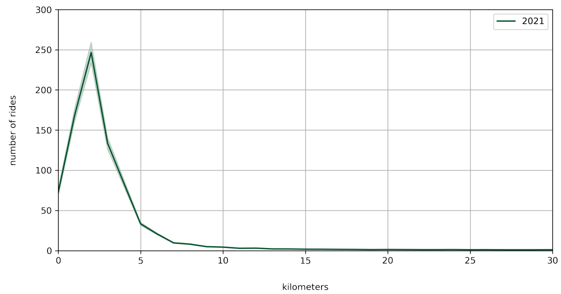

(a) Traveled distance

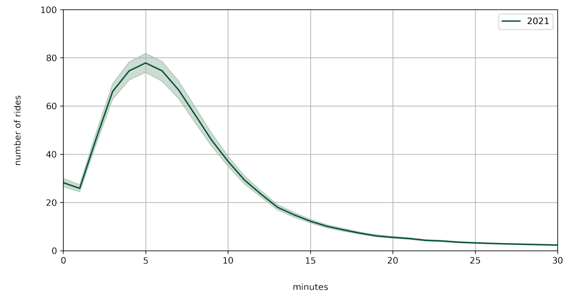

(b) Duration

Figure 13: Distribution of daily traveled distance and duration of rides in Eindhoven with a 95% confidence interval.

After completing a ride, the user leaves the moped at its destination. The moped will then be available to other users to make a ride. Whether or when the moped is used again for a next ride depends on the demand at this time and location. The elapsed time between the end of the ride and the start of a new ride for the same moped is referred to as the idle time. The idle time is linked to the end location of the first ride. Figure 14 depicts the average daily idle time throughout the year. The pattern in this case is roughly the inverse of the pattern associated with the number of daily rides in Figure 7. The idle times are lower in the summer months and higher in the winter months. This makes sense because when more rides take place, the mopeds will spend less time idle.

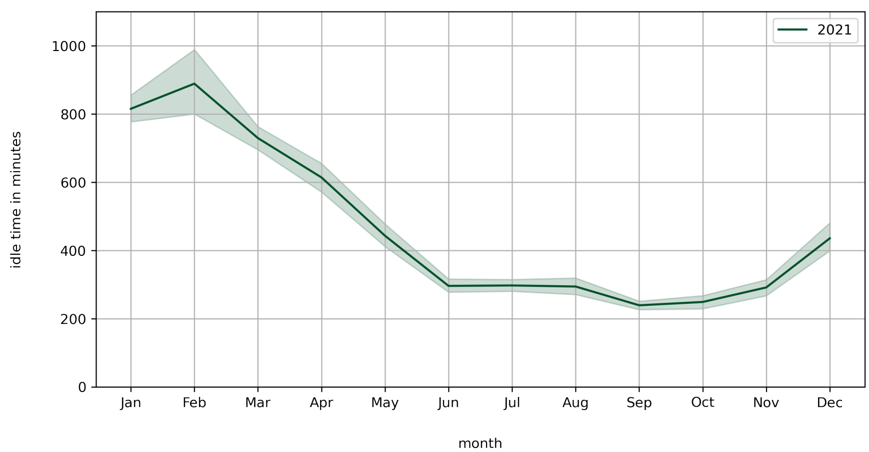

Figure 14: Average daily idle time in Eindhoven with a 95% confidence interval.

If the idle time is presented by the day of the week, rather than the month of the year, an interesting insight emerges, shown in Figure 15. Here, it is evident that the average daily idle time decreases as the weekend approaches, reaching a low on Friday. The lowest point here on Friday does not correspond to the highest point on Saturday from Figure 7, indicating that there is no direct one-to-one correlation.

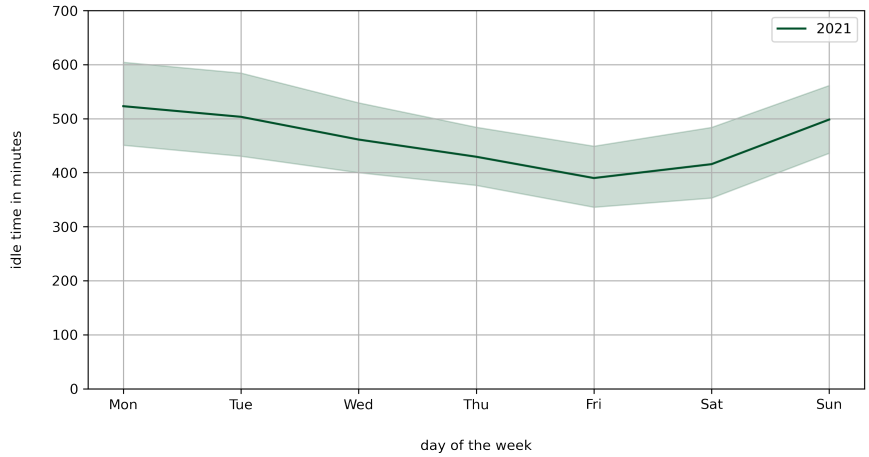

Figure 15: Average daily idle time in Eindhoven with a 95% confidence interval.

Aside from the fact that the average daily idle time vary by the month of the year and the day of the week, they also vary throughout the day. Figure 16 depicts the average idle time per hour of the day, segmented by the day of the week. It can be observed here that the lowest idle times are around noon, while the highest are around midnight. This seems plausible, given the low utilization of the mopeds during these night hours. In addition, the pattern of the weekdays does not much differ from the weekend days, which was observed at the average number of hourly rides in Figure 9.

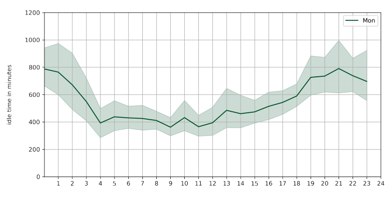

(a) Monday

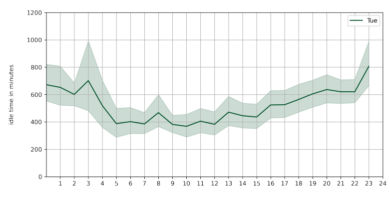

(b) Tuesday

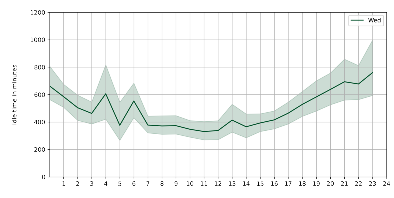

(c) Wednesday

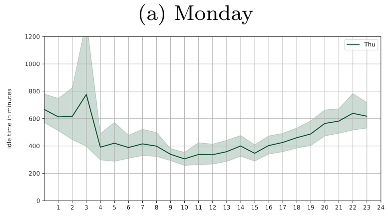

(d) Thursday

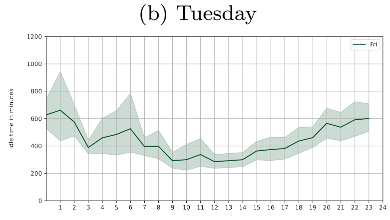

(e) Friday

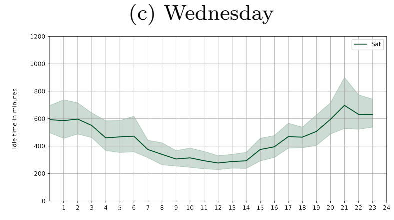

(f) Saturday

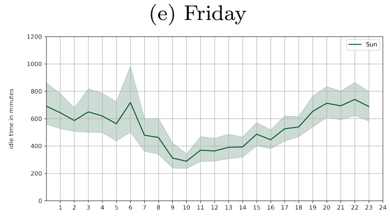

(g) Sunday

Figure 16: Average hourly idle times in Eindhoven with a 95% confidence interval.

With the figures above, it is easy to see that the idle times vary depending on the hour of the day, but they might also differ substantially depending on location. Therefore, it is wise to also look at the variations in idle times in a network. A grid system is frequently used to make data clearly visible in a network, with each grid reflecting the data of the points beneath it. To work with geospatial data, Felyx uses H3, an open-source indexing system with a hexagonal grid developed by Uber [75]. This framework enables any location on the planet to be assigned to different-sized hexagons, each with its own unique ID.

Figure 17 depicts a 2D and 3D visualization of the average idle time for each hexagon. The colors represent the idle times, where yellow denotes the lowest numbers, followed by orange, red, and finally purple with the highest values. It clearly shows that the idle times are lower in the city center and higher on the outskirts of the service area.

Figure 17: A 2D and 3D visualization of the average idle time on a hexagon level for Eindhoven in 2021.

4.1.2 Summary

Since its start, Felyx has grown significantly. The fleet size increased, the mopeds are available in more cities, and Felyx is much more well-known among its audience. This is reflected in the data, such as the increase in rides year over year. This expansion has a strong seasonal component, resulting in higher use in the summer and lower use in the winter. In addition, the type of weather plays a major role in the utilization of the mopeds. It is decided to further investigate the data of one city, which is Eindhoven.

The number of daily rides that take place is time-dependent. This indicates that the number varies by the month of the year, the day of the week, and the hour of the day. The opposite of the number of rides is the time that the mopeds are not used, commonly referred to as the idle time. According to the data, the average idle time is significantly lower when the mopeds are used frequently compared to when they are used less frequently, resulting in higher idle times. The average idle time also varies, just like the number of rides that take place, by the month of the year, the day of the week, and the hour of the day. Additionally, the location of the moped has a major impact on the idle time. Mopeds that are located more in the center of the city tend to have a much lower idle time than mopeds located on the outskirts of the city center. Worth mentioning here is that there is a pattern in the number of daily rides during weekdays that differs from the pattern on weekend days. This could be due to people who use a Felyx moped to commute throughout the week, while others have different purposes during the weekends. However, this pattern cannot be identified in the average idle time per day of the week.

The data analysis also revealed that there is a distribution in the traveled distance and duration per ride. These two factors depend of course on the start and end location of the rides, but they are also user dependent and may be affected by traffic at that time of the day.

Furthermore, it has been noted that certain areas of the city are more frequently used as the start and/or end points of a ride. This is illustrated with heat maps which show that the transitions, where rides start and end, vary depending on the day of the week, and the hour of the day.

To conclude, this data analysis has taught us that because the data varies significantly depending on the month of the year, the day of the week, and the hour of the day, but also on the location, it should also be used in that manner to predict future occurrences. Fortunately, the data is not completely random, and patterns are recognized, allowing some data to be combined to provide more precise predictions. In other words, in general, when predicting what will happen on a Monday morning at a certain location, only data from Monday mornings at this location will be considered.

4.2 Simulation

The real-world system of shared mobility operator Felyx is modeled using the discrete-event simulation model proposed in Section 3.2.1. By using this simulation model, ultimately, the impact of various system modifications can be investigated without implementing them in the real world, which in most circumstances saves both time and money. It is important to keep in mind that simulation models generally represent reality in a simplified manner, which might lead to some discrepancies in the results when compared to real data.

In the simulation model, distributions are used for the idle times, the transitions, and the travel times. It is therefore necessary to execute multiple simulation replications since every simulation run will produce different results. Therefore, an analysis is first carried out to determine the number of simulation replications required to obtain accurate results within a reasonable amount of time. The behavior of the sample mean as the number of replications increases is examined for the four KPIs (Section 3.2.1). This analysis is presented in Figure 18 for two randomly selected days in Eindhoven. Here, it can be seen that the sample mean for all four KPIs starts off more erratic but stabilizes and converges as the number of replications increases. Considering time and accuracy, it is decided to use 50 simulation replications in the remainder of this study.

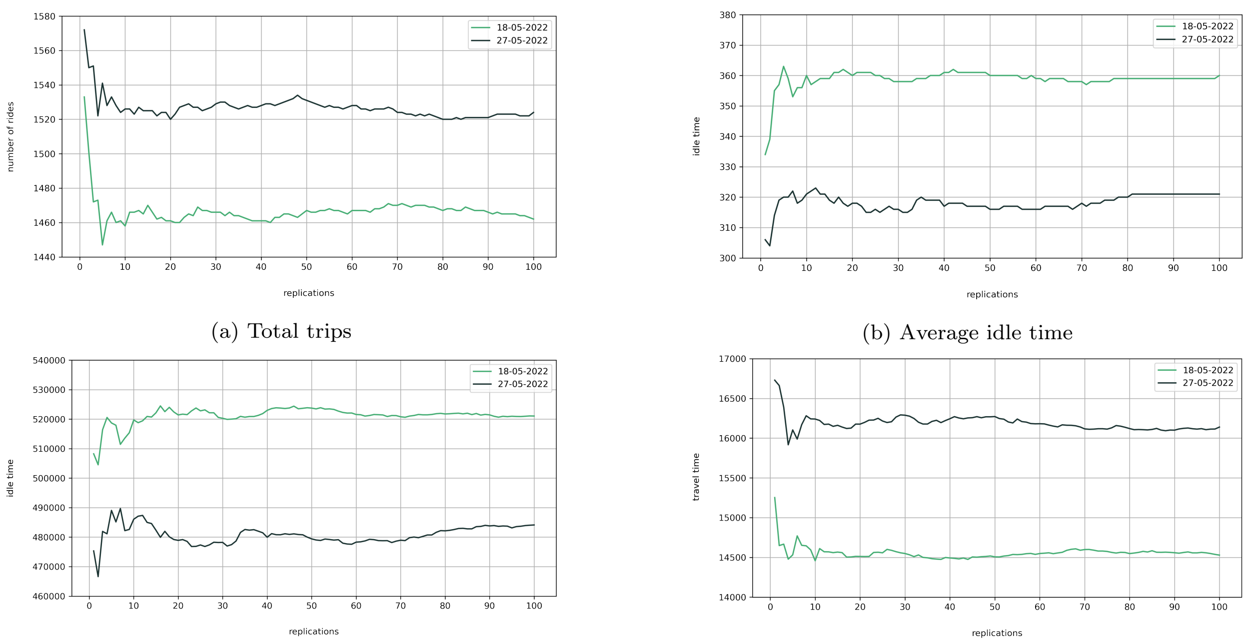

(a) Total trips (b) Average idle time

(c) Total idle time (d) Total travel time

Figure 18: The sample mean of the KPIs with an increasing number of replications.

Several simulation model variants are evaluated and compared to real data. First of all, various granularity levels of the H3 framework [72] are tested. Felyx primarily uses level 9 for spatial analyses; however, to include more data points per hexagon and hence improve the accuracy of the distributions, level 8 and 7 are also examined. Figure 19 depicts the service area of Eindhoven, divided into hexagons of level 9, 8, and 7 respectively. The hexagons' side lengths are approximately 200, 500, and 1400 meters, accordingly.

(a) Level 9 (b) Level 8 (c) Level 7

Figure 19: The service area of Eindhoven divided into various hexagon granularity levels.

Secondly, in some variants of the simulation model, the type of weather is incorporated. It could well be that the idle times and the transitions are affected by whether it's a sunny or rainy day. Finally, experiments are carried out to assess whether the results of the simulation model would improve if the idle times, which are currently assumed to depend only on the day of the week, were considered to differ per hour of the day.

4.2.1 Results

The simulation model is evaluated on Eindhoven considering one month of data, starting from the 1^{st} to the 31^{st} of May 2022. According to all the above-mentioned variants and experiments, the simulation model performs best when granularity level 8 is used, the weather type is not included, and the idle times only depend on the day of the week. The comparison of these results to real data is explained below. The time required to simulate one replication of a day is around 5 seconds on a desktop with 6-Core Intel Core i7 CPU 2.6 GHz processor, 16GB memory, and macOS 12.6 Monterey.

To assess the simulation model's performance, the simulated data is analyzed at several stages. First of all, it is compared to real data using the KPIs, as can be seen in Figure 20. The four graphs depict both the real data and the simulated data per day for the entire month. The sample means and accompanying standard deviations from 50 simulation replications are used to represent the simulated data. It is clear that the simulated data does not exactly match the real data. However, because a simulation model is a simplified representation of reality, this will never be the case. To assess the simulation model's performance based on the KPIs, the mean absolute percentage error (MAPE) is determined. The results for each KPI are, respectively: 19.21%, 11.69%, 16.23%, and 30.54%, which leads to an overall MAPE of 19.42% on average. The results of the KPIs for the comparison of the other variants of the simulation model to real data can be found in the Appendix.

The above mentioned simulation's results are discussed with people from Felyx. They considered the differences between the simulated data and the real data to be reasonable, given the stochasticity present in such a context, something they also observe in other prediction models.

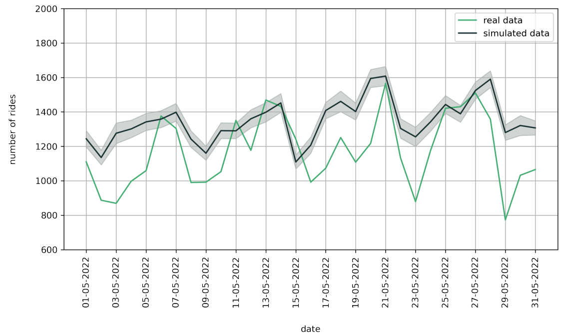

(a) Total trips

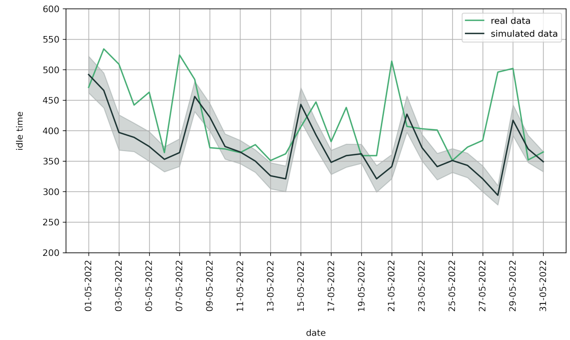

(b) Average idle time

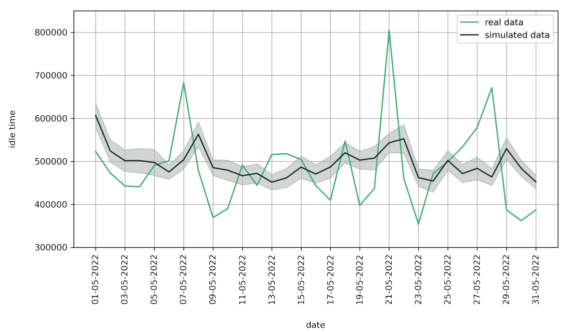

(c) Total idle time

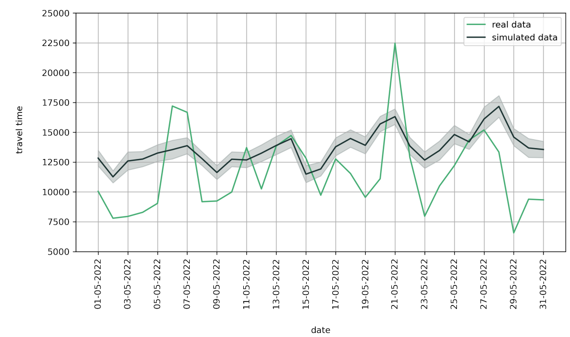

(d) Total travel time

Figure 20: Real and simulated data of the KPIs.

In addition to the KPIs, the start location of the rides in the simulation are compared to the start location of the rides in reality. The end location of the rides is not included here because the end location of a previous ride is the start location of the next ride. A detailed comparison of the start location of these rides is carried out on a randomly selected day, the 14$^{\text{th}}$ of May, which is visualized in Figure 21. This figure illustrates all hexagons within the service area for both the real and simulated data. The numbers in the hexagons represent the percentage of the total number of rides that started on this day for each particular hexagon. The percentages are rounded to the nearest integer. As can be seen in both images, the majority of the rides started in the hexagon where the central station of Eindhoven is located, which is indicated as a dark green hexagon. On this specific day, over the whole service area, there is an average absolute difference of 0.39 per hexagon between the real and simulated data. Furthermore, the average deviation within the 10 hexagons with the most rides is 17.51%.

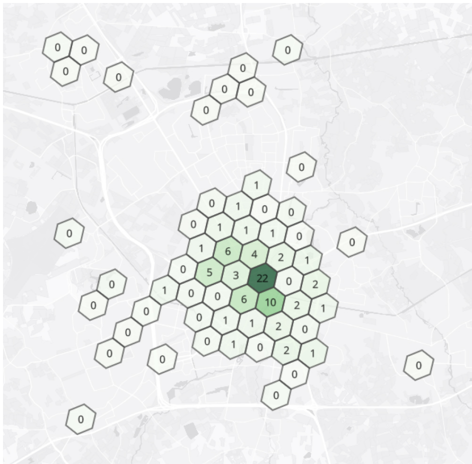

(a) Real data

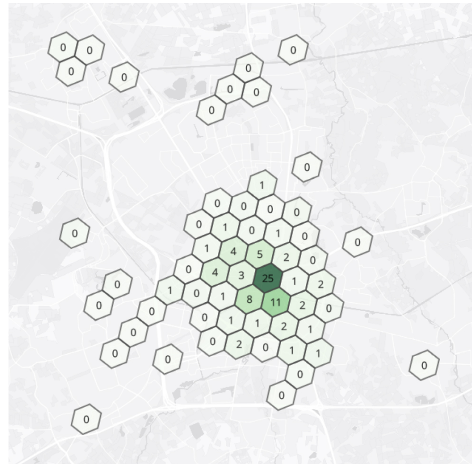

(b) Simulated data

Figure 21: The percentages of the total number of rides per hexagon on May 14 for real and simulated data.

The comparison for this specific day is also conducted for the whole month of May. The results are depicted in Figure 22. As can be observed, on some days the simulated data matches the real data more closely than others. Over the entire month, there is an average absolute difference of 0.46 per hexagon between the real and simulated data, followed by an average deviation for the 10 hexagons with the most rides of 24.30%.

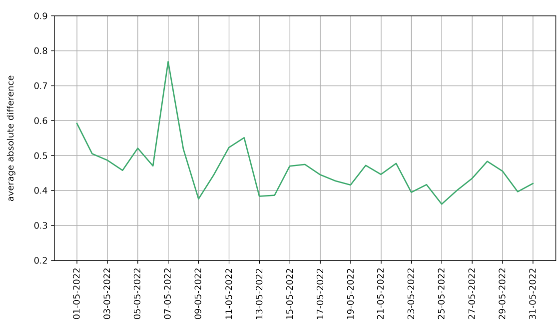

(a) Average absolute difference

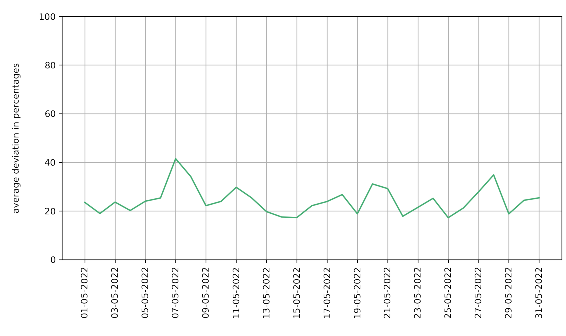

(b) Average deviation

Figure 22: Average absolute difference and deviation per hexagon between real and simulated data.

4.3 Simulation-based optimization

Now that the performance of the simulation model is established, the following step is the integration with the optimization model proposed in Section 3.2.2. This optimization model makes use of information about the simulation's events and is therefore able to find the best actions to take to potentially improve the system.

In this study, the actions to improve the system consist of rebalancing the vehicles, which in Felyx's case are mopeds. The initial distribution of the fleet of mopeds across the service area can be modified to affect and hopefully enhance the simulation's outcomes. Because Felyx eventually intends to increase revenue, the results of the KPIs mentioned in earlier chapters are converted into a revenue calculation. This is done in the following way: Users of the Felyx mopeds pay a one-time fee of 0.50 euros to start a ride, followed by a fee of 0.30 euros per minute. To calculate the revenue, the number of rides made during the simulation period is multiplied by 0.50 euros and then added to the total minutes driven multiplied by 0.30 euros. The minutes driven are rounded up for each ride, as is also the case in reality.

4.3.1 Results

Two different scenarios are evaluated using two variants of the simulation-based optimization model. In scenario 1, vehicles are rebalanced from locations with high idle times to locations with low idle times. In scenario 2, this is also the case but this time it is also considered whether there is a large surplus or potential deficit of vehicles at the pickup and drop-off locations. For both scenario 1 and 2, the impact of rebalancing actions on the generated revenue is investigated throughout the month of May. The effect of rebalancing 5, 10, 20, and 50 mopeds at once is examined.

Scenario 1

The simulation-based optimization model for scenario 1 determines each time what the optimal rebalancing action is and then changes the initial distribution of the mopeds to this new situation. The results should indicate that a moped must be picked up from the location with the highest average idle time and the presence of at least one moped. As a result, the number of available mopeds at this location decreases by one. The moped that is just picked up must be dropped off at the location with the lowest average idle time, where the number of mopeds will increase by one. When determining the next rebalancing action, it might be that no other moped can be picked up from the location a moped was just picked up, because there are none available anymore. In that case, the next location with the highest idle time and the presence of at least one moped is considered. For the first of March, the rebalancing actions with the various rebalancing capacities are visualized in Figure 23.

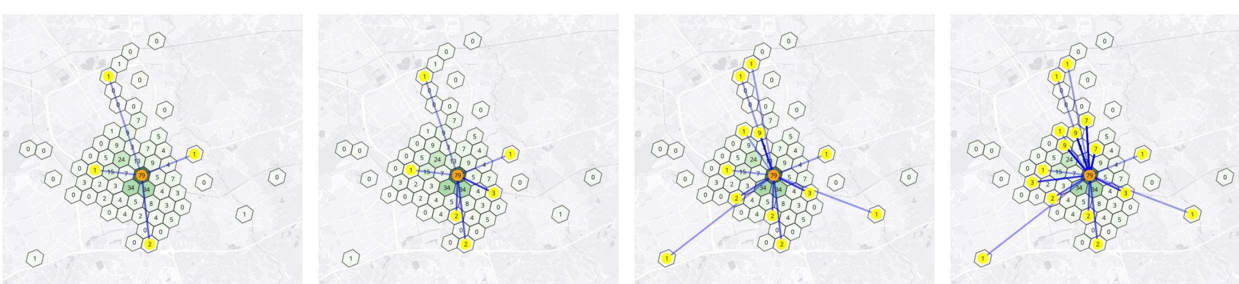

(a) rc = 5 (b) rc = 10 (c) rc = 20 (d) rc = 50

Figure 23: Rebalancing actions on May 1 for scenario 1 with various rebalancing capacities (rc).

The hexagons indicate the locations that are used throughout the simulation replications, whereas the numbers represent the initial distribution of the mopeds. The mopeds must be picked up from the hexagons with yellow marks and dropped off at the hexagons with orange marks. The opacity of the blue line reflects the number of mopeds that must be rebalanced. The higher the opacity, the more mopeds to rebalance. To view this more clearly, the rebalancing actions will, of course, also be listed in tables.

An important thing that stands out in Figure 23, which was expected, is that for every rebalancing capacity, all mopeds must be dropped off at the same hexagon. The lowest idle time is always found there, which explains why. However, this won't always be a realistic situation because, at some point, there could be an oversupply of mopeds at this hexagon which definitely decreases the service level to the users. The number of mopeds then exceeds the demand, resulting in many mopeds standing idle. The reason for this is that in the simulation, the idle time distribution for each hexagon is based on idle times from historical data for this particular hexagon and day of the week. The number of mopeds that was present at the times of the data points of the idle time distributions is unknown. It can be said that the idle time distributions are based on the average number of mopeds belonging to the hexagons. This means that the idle time distributions are not directly linked to the number of mopeds in these hexagons and will therefore not change as more or fewer mopeds are present. Although, there's a high probability that in reality at a location with a large surplus of mopeds, the idle time per moped will definitely increase. This would be the case for the hexagon with the orange mark where an additional 5, 10, 20, or even 50 mopeds are added on top of the initial number of mopeds. A more realistic situation where this is prevented is in scenario 2.

Scenario 2

The simulation-based optimization model for scenario 2 again determines each time what the optimal rebalancing action is and then changes the initial distribution of the mopeds to this new situation. This time, the results should indicate that not all mopeds must be dropped off at the same hexagon, but that it depends on where there are potential deficits of mopeds. If there appears to be a potential deficit in any hexagon and a moped is dropped off here, the deficit decreases by one. This may resolve the deficit of mopeds at this hexagon, implying that the following mopeds must be rebalanced to other hexagons. Figure 24 visualizes, again for the first of March, the rebalancing actions with the various rebalancing capacities. It is clearly visible here that the drop-off locations vary from the drop-off locations in scenario 1.

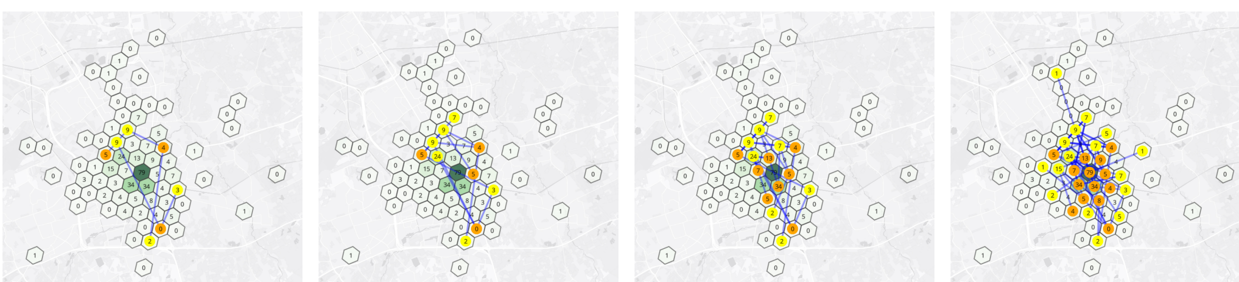

(a) rc = 5 (b) rc = 10 (c) rc = 20 (d) rc = 50

Figure 24: Rebalancing actions on May 1 for scenario 2 with various rebalancing capacities (rc).

Now that the optimal rebalancing actions to perform at the start of the simulation are known, a comparison can be made between the system with the initial distribution of the mopeds and the system with the distribution of the mopeds after the rebalancing actions. In both systems, the locations of the mopeds are used as the starting point of the simulation. After running the simulation, the total number of trips and the total travel time of all mopeds are then used to determine the generated revenue per system (see Section 4.3). Figure 25 depicts the total trips, total travel time, and revenue increase in percentages for the month of May after performing the rebalancing actions with various rebalancing capacities.

Figure 25: Total trips, total travel time, and revenue increase in percentages for scenario 2.

It can be seen that there is an increase in the total trips and total travel time for all the rebalancing capacities on almost every day of the month, leading to an increase in revenue in 98.4% of the situations. It is also clear that the quantity of the revenue increase fluctuates depending on the day and the rebalancing capacity. On some days, performing rebalancing actions lead to a larger increase in revenue than on others. A closer investigation is conducted on two days that show contrasting behavior, specifically May 17, where there is hardly any revenue increase possible, and May 29, which shows a lot of potential.

Figure 26 depicts the rebalancing actions for both days with a rebalancing capacity of 20 mopeds. By examining the initial distribution of the mopeds, which is represented by the numbers in the hexagons, it can be observed that on May 17, the mopeds are primarily located in the city center. As a result, rebalancing some mopeds to nearby locations in the city center does not immediately result in extra revenue. On May 29, however, there is a significant number of mopeds located outside of the city center. By rebalancing these mopeds in particular, extra revenue can easily be generated.

(a) May 17

(b) May 29

Figure 26: Rebalancing actions on two days that show contrasting behavior with rc = 20.

Although the simulations for the system with the initial distribution of the mopeds and the simulations for the system with the distribution of the mopeds after performing the rebalancing actions are both conducted for a set of 50 replications each day, there will always be a slight difference in the results if again a set of 50 replications is conducted for the same day. This is because each run of the simulation produces a unique result since there is sampled from continuous distributions. By conducting multiple replications, the sample mean will approach the simulation's true value.

If conducting a set of 50 replications per system per day is repeated for, in this case, five times, a total of five sample means is obtained per system. Subsequently, the mean of these five sample means can be determined as well as the standard deviation of these sample means. This standard deviation of the mean of the sample means is also known as the standard error (SE). The SE tells how much the sample mean would vary if a study is repeated with new samples from the same population. In this case, if another set of 50 replications per system per day is conducted. The mean of the sample means as well as the standard errors are shown in Figure 27. It can be observed that the graphs occasionally overlap. In these situations, it cannot be stated with such confidence that performing the rebalancing actions leads to an increase in revenue.

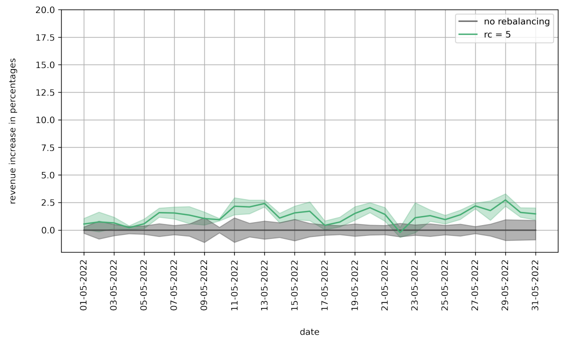

(a) rc = 5

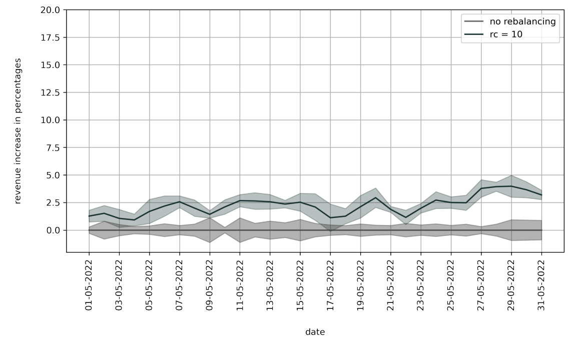

(b) rc = 10

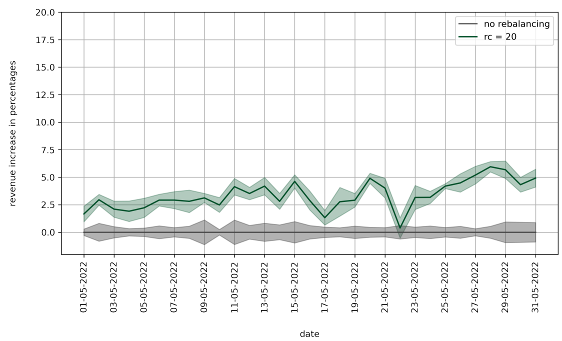

(c) rc = 20

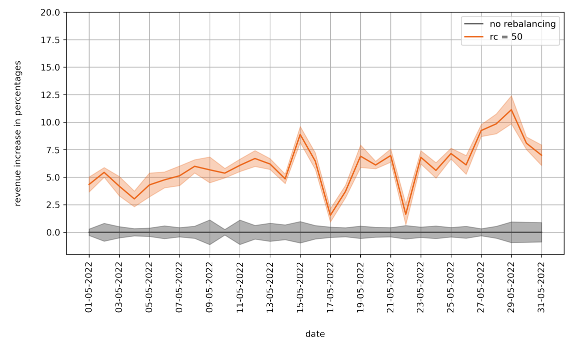

(d) rc = 50

Figure 27: Revenue increase in percentages with standard errors (SE) for scenario 2.

An important note to emphasize is that another study by Felyx revealed that rebalancing actions affect the generated revenue not just on the day they are performed, but also on the days following. Mopeds that may have been idle for several days at their initial locations are now rebalanced to locations with higher demand. From these new locations, users take the mopeds and drive to other areas, where demand is again likely to be higher than at its initial location, resulting in several extra rides over multiple days. According to this study, the effect can be noticed for up to three days.

In the results of Figure 25 and Figure 27, the model is not used as a continuum. On each subsequent day, the model retrieves the locations of the mopeds from real data and uses these as the starting point for the simulation on that day. As a result, the distribution of the mopeds at the end of the previous day differs from the distribution of the vehicles at the start of the next day. However, to determine the effect of rebalancing actions over several days, the model must be used as a continuum. Therefore, the model is adjusted and now keeps track of the distribution of the mopeds at the end of the day and uses it as the starting point of the simulation on the following day. It is assumed that there's no activity during the night because only 3.48% of the rides in 2021 occurred during these hours. The model now only requires the locations of the mopeds from real data at the start of the simulation on the first day and may thereafter run over several days. The following days can therefore also lie in the future.

Figure 28 depicts the revenue increase in percentages over three consecutive days, where the rebalancing actions are only performed on the first day, which are in this case May 13 and May 27. The error bars indicate the standard error (SE) which is calculated from five model replications. It can be seen that there is an increase in revenue not only on the first day but also on the two following days. For the 13th of May, the total revenue increase over three days is on average 1.25%, 2.05%, 3.22%, and 4.35% for rebalancing 5, 10, 20, or 50 mopeds, respectively. For the 27th of May, this is 1.59%, 2.36%, 3.77%, and 6.57%.

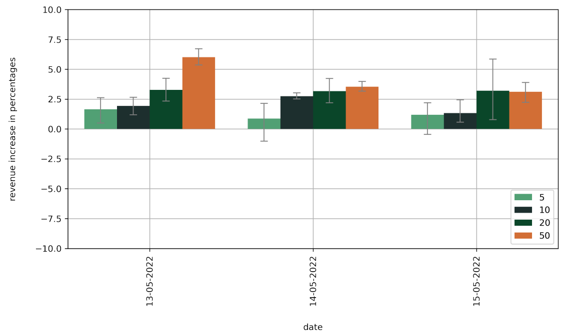

(a) May 13 - May 15

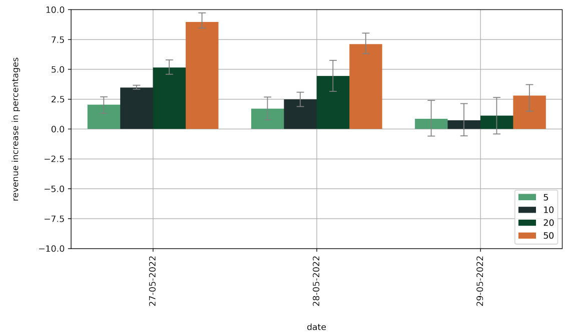

(b) May 27 - May 29

Figure 28: Revenue increase over three consecutive days due to rebalancing actions on the first day.

4.3.2 Costs of rebalancing

In order to asses whether the extra revenue generated by the rebalancing actions also results in extra profit, several costs must be considered. Based on conversations with multiple Felyx employees, the process of performing the rebalancing actions and the related costs are drawn out.

Felyx performs the rebalancing actions by using a van that can hold up to five mopeds. This van is driven by an employee who performs the rebalancing actions on his own. During an 8-hour shift, an employee is able to rebalance 25 mopeds on average, which is equivalent to rebalancing five mopeds in 96 minutes. The average monthly cost of a van, including fuel, insurance, and maintenance, is 1000 euros. Assuming that this van is used seven days a week for eight hours a day, the van cost 0.07 euros per minute. The cost for an employee is 20 euros per hour, which is equal to 0.33 euros per minute. The total costs for the various rebalancing capacities are shown in Table 5. It is worth noting that this is a simplified assumption for the costs. In general, a van or an employee is not hired per minute. However, the employees at Felyx also perform other tasks during a working shift, such as repairing mopeds. Also, the van is used for other purposes, such as replacing the batteries of the mopeds. As a result, it is decided to work with a cost per minute, resulting in a linear distribution.

Table 5: Rebalancing costs for the various rebalancing capacities.

Rebalancing capacity	Total costs
5	€ 38.40
10	€ 76.80
20	€ 153.60
50	€ 384.00

Additionally, it must be taken into account that the extra revenue is a result of more rides and a larger total travel time of the fleet. Due to this increased use, there is slightly more depreciation of the mopeds, the batteries need to be recharged more frequently, and there is a little higher risk of damage. As a general guideline, Felyx employs a percentage of 30%, which must be subtracted from the extra revenue to cover these costs.

Profit

When the costs are taken into consideration, the final extra profit for the same three consecutive days as described before is as follows (see Figure 29). A comparison of the revenue and profit graphs reveals that on the first day, the profit is significantly lower, if not negative, than the revenue. This is due to the costs of the rebalancing operations, which only take place on the first day. Additionally, 30% of the revenue is withheld in all cases to cover some general operational costs, as described in Section 4.3.2. For the 13th of May, the total profit increase over three days is on average 0.58%, 0.84%, 1.07%, and 0.10% for rebalancing 5, 10, 20, or 50 mopeds, respectively. For the 27th of May, this is 0.85%, 1.14%, 1.62%, and 2.06%.

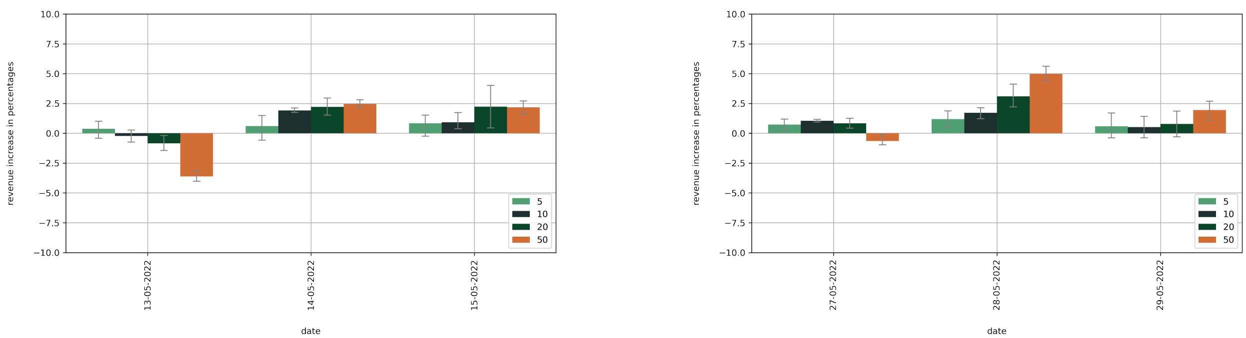

(a) May 13 - May 15

(b) May 27 - May 29

Figure 29: Profit increase over three consecutive days due to rebalancing actions on the first day.

Felyx rebalances the vehicles in their fleet on a regular basis, on average about once every three days during low season and once every two days or even every day during high season. The profit growth in euros can therefore rise significantly over the course of a year. Additionally, Felyx performs rebalancing actions in each city in which they operate, raising this number even higher.

5 Conclusions and future research

To conclude this study and answer the main research question, first, the sub-questions are addressed. While addressing these questions, the main findings are shared along with any limitations. Finally, suggestions for future research are given.

The results of the data analysis on the real data of the shared moped provider Felyx reveal that there are definitely patterns to be found. In the Netherlands in general, there is a strong seasonality in the usage of the mopeds. A plausible explanation for this is that people tend to go out more in the summer since the weather is nicer. When analyzing a specific city, in this case, Eindhoven, this seasonality can also be seen. Examining the data by the day of the week shows that there is an increase in moped use as the weekend approaches. During weekdays, this usage follows a similar pattern with a morning, afternoon, and evening peak. The weekend days, on the other hand, exhibit a different pattern. What the data analysis even more reveals is the following. First, the idle times, which is the time between two rides that a moped is not used, vary throughout the year, as well as by the day of the week and the hour of the day. Secondly, the transitions, or where rides go, do not differ much throughout the year, but rather by the day of the week and the hour of the day. Lastly, the travel times hardly differ throughout the year, regardless of the day of the week or the hour of the day. One thing all three components have in common: they are location dependent.

With this information, a discrete event simulation model is created that imitates the behavior of the free-floating shared mobility system of Felyx by using a hexagonal grid system. The simulator uses historical data up to the simulation's starting point as input and can therefore simulate periods from the past as well as periods from present time into the future. It runs for a sufficient number of replications and calculates the sample means and standard deviations of four KPIs, which are then compared to real data. Experiments with several variants of the simulation model are conducted to optimize its performance. The simulator is evaluated on the whole month of May '22, where the results show an average MAPE of 19.42% based on the four KPIs. In addition to the KPIs, the simulator is also evaluated based on the locations visited. The results show an average absolute difference of 0.46 per hexagon and an average deviation for the 10 hexagons with the most rides of 24.30% between the simulated data and the real data. The simulation's results are discussed with people from Felyx. They considered the differences between the simulated data and the real data to be reasonable, given the stochasticity present in such a context, something they also observe in other prediction models.

Following that, the simulation results are used in an optimization model. First, the average idle time in the simulation is assessed per location. With this, the optimal rebalancing actions are determined which indicate that all vehicles must be rebalanced to the location with the lowest average idle time in order to maximize the KPIs of the simulation. This is not a realistic scenario, as rebalancing many vehicles to this location would result in a surplus of vehicles at this location. For that reason, secondly, also the surpluses and potential deficits throughout the simulation are determined per location. The optimal rebalancing actions then indicate that the vehicles must be rebalanced from locations with a high average idle time and a surplus of vehicles to locations with a low average idle time and a potential deficit of vehicles.

With the optimal rebalancing actions established, the results of the simulation of the system with the initial distribution of the vehicles can be compared with the results of the simulation of the system with the distribution of the vehicles after the rebalancing actions have been performed. In the case study in this research, this comparison is made for the month of May '22 and shows an increase in the total trips and total travel time, which in addition to an increase in service level also leads to an increase in revenue in 98.4% of the situations. This increase in revenue can reach up to 11.13% by rebalancing 50 mopeds on a specific day. The comparison also reveals that the impact of rebalancing varies per day, which is mainly due to the imbalance between supply and demand being greater on certain days than others. Crucial here is that rebalancing actions affect the generated revenue not just on the day they are performed, but also on the days following. Therefore, the total extra revenue that is generated over three days, where rebalancing actions are only performed on the first day, is examined. The results from two distinct days demonstrate an increase in revenue over three days of 4.35% and 6.57% by rebalancing 50 mopeds. In order to calculate the extra profit, the costs of the rebalancing operations must, of course, be taken into account. These costs vary depending on the rebalancing capacity. The final results show a possible increase in profit of 1.07% and 2.06% for the two days that are examined. Rebalancing actions usually take place multiple times per week in every city an operator is active in and can therefore lead to significant profit growth in euros over the course of a year.

Limitations

Aside from the fact that certain patterns can be recognized in the real data of shared moped provider Felyx, it remains difficult to predict how the system will behave. Felyx is currently a scale-up with strong yearly growth, which results in inconsistent changes in the data.

In the simulation model, the idle times are eventually not defined by the hour of the day, but merely by the day of the week, while the data analysis revealed that the idle times do vary throughout the day. The reason for this is when specifying the idle times per hour, the number of data points decreases, making it more difficult to construct suitable distributions out of these data points. To account for this problem, the current data points can be extrapolated, however, incorporating the idle times per hour did not improve the results based on the KPIs. Additionally, the data analysis also revealed that the type of weather can have an impact on the usage of the vehicles. However, the results of the variant of the simulation model where the weather type is included also showed no improvement. Again, by further specifying the data, the number of data points decreases, which might lead to increased stochasticity. Thirdly, the idle times have a significant impact on the simulation's results. A correct distribution of the idle times per location is difficult to obtain, usually, the idle times are also affected by the number of vehicles present. This variation in the idle times due to more or less vehicles present is not accounted for in the simulation model. Finally, keep in mind that the simulation model is a simplified representation of reality. There are still many factors that are not accounted for in the simulation that do influence the usage of the vehicles. To give a few examples, consider that the batteries of the vehicles run out or a vehicle breaks down, making it (temporarily) unusable. In addition, there might also be some errors in the historical data that the simulation uses, such as GPS sensors that have reported incorrect positions.

The results obtained with the simulation-based optimization model will vary significantly depending on the shared mobility system under consideration, as well as on the city, the season, and even the day. It is therefore challenging to assess the model's real performance. To be more explicit, if the use of shared vehicles in a shared mobility system is higher in a certain period, performing rebalancing actions will result in extra revenue than if the vehicles are used less frequently. The costs of rebalancing operations, on the other hand, will remain constant, resulting in more profit. If the use of the vehicles is low, it may be required to perform rebalancing operations to better match the demand (or increase the service level), which may result in a loss at the moment, but will eventually lead to a profitable operation if the use of the vehicles rises.

Future research

The simulation-based model currently calculates the system improvement if 5, 10, 20, or 50 vehicles are rebalanced. In some cases, however, performing 50 rebalancing actions may have almost the same results as performing 30 rebalancing actions. The tiny additional benefit between performing 30 to 50 rebalancing actions may not be worth the costs. The model can be adjusted in such a way that it stops rebalancing as soon as the improvement yields less than a certain threshold. Another approach could be to examine the surpluses and potential deficits in the system, which are currently already included in the optimization model, and that the model stops rebalancing once this imbalance is reduced to a certain value.

It is recommended to perform the rebalancing actions from the simulation-based optimization model using a vehicle routing algorithm. To be specific, performing the rebalancing actions can be seen as a capacitated vehicle routing problem with pickup and delivery (CVRPPD). Using this algorithm, the optimal route can be determined to perform the rebalancing actions, taking into account the capacity of the rebalancing vehicle. First of all, this algorithm is also able to better draw out the actual costs, and secondly, various rebalancing strategies may be assessed. Consider, for example, the use of different vehicles with varying capacities.

Finally, combining operator-based rebalancing with user-based rebalancing could be an interesting subject for future research. User-based rebalancing can sometimes be more effective due to its low costs. However, this alone is insufficient, which is why a combination with operator-based rebalancing is required. The development of a single algorithm in which both types of rebalancing are combined might result in an even greater improvement in service quality as well as in generated profit, which is worth investigating.

6 References

[1] S. Shaheen, N. Chan, "Mobility and the Sharing Economy: Potential to Overcome First- and Last-Mile Public Transit Connections", *Research Papers in Economics* **2016**.

[2] L. Davis, "Rolling along the last mile: Bike-sharing programs blossom nationwide", **2014**.

[3] J. Ploeger, R. Oldenziel, "The sociotechnical roots of smart mobility: Bike sharing since 1965", *The Jounal of Transport History* **2020**.

[4] The Guardian, How this Amsterdam inventor gave bike-sharing to the world, https://www.theguardian.com/cities/2016/apr/26/story-cities-amsterdam-bike-share-scheme, **2016**.

[5] Uber, About us, https://www.uber.com/us/en/about/.

[6] John Michel, Why Uber is Poised to Dominate the MaaS Market, https://medium.com/@john.michel1927/why-uber-is-poised-to-dominate-the-maas-market-b581bd6ca2f9, **2018**.

[7] Ryan Lawler, "Zipcar For Scooters" Startup Scoot Networks Launches To The Public In San Francisco, https://techcrunch.com/2012/09/26/scoot-sf-launch/, **2012**.

[8] Abel de Graaff, Hoe twee 29-jarigen met elektrische deelscooters Amsterdam een beetje schoner willen maken, https://www.businessinsider.nl/hoe-felyx-met-elektrische-deelscooters-amsterdam-een-beetje-schoner-wil-maken/, **2017**.

[9] Megan Rose Dickey, The electric scooter wars of 2018, https://techcrunch.com/2018/12/23/the-electric-scooter-wars-of-2018/?guccounter=1, **2018**.

[10] Mattia Di Francesco, Is the Urban Air Mobility Industry ready for Take-Off?, https://www.airborne.com/urban-air-mobility-the-rise-of-evtol-vehicles/.

[11] Government of the Netherlands, Number of electric shared cars has doubled in 2021, https://www.government.nl/latest/news/2021/11/04/number-of-electric-shared-cars-has-doubled-in-2021, **2021**.

[12] S. Shaheen, A. Cohen, "Mobility on demand (MOD) and mobility as a service (MaaS): early understanding of shared mobility impacts and public transit partnerships", *Demand for Emerging Transportation Systems* **2020**.

[13] A. C. Change, E. P. (ACCEPT), Carpooling for the climate, https://accept.aseanenergy.org/carpooling-for-the-climate/, **2020**.

[14] A. Tips, Car Sharing in Amsterdam, https://www.amsterdamtips.com/car-sharing-amsterdam, **2021**.

[15] Zentralplus, Luzerner Regierung will von Uber-Verbot nichts wissen, https://www.zentralplus.ch/politik/luzerner-regierung-will-von-uber-verbot-nichts-wissen-1813391/, **2020**.

[16] Z. Jiangping, Z. Wadud, J. Jiao, Y. Wang, "Understanding and planning shared micro-mobility", *Transportation Research Part D: Transport and Environment* **2022**.

[17] R. Basselier, G. Langenus, L. Walravens, "The rise of the sharing economy", *Economic Review* **2018**.

[18] T. Yoon, C. R. Cherry, L. R. Jones, "One-way and round-trip carsharing: A stated preference experiment in Beijing", *Transportation Research Part D: Transport and Environment* **2017**.

[19] F. Palmieri, VAIMOO al MOVE 2021 presenta "LIBERA" l'e-bike connessa, https://www.ceotech.it/vaimoo-al-move-2021-presenta-libera-le-bike-connessa/, **2021**.

[20] R. Smit, Delftse mysteries: waar zijn al die mobikes gebleven?, https://indebuurt.nl/delft/genieten-van/mysteries/delftse-mysteries-waar-zijn-al-die-mobikes-gebleven~103390/.

[21] N. Trautmann, M. Gnägi, *Operations Research Proceedings 2021*, Springer Publishing, **2022**.

[22] D. Cataneo, The smart way to balance supply and demand in a shared micromobility, https://www.electricfeel.com/blog/supply-demand-shared-micromobility, **2022**.

[23] E. Fields, Supply-demand imbalance in shared mobility: cost and consequence, https://medium.com/zoba-blog/supply-demand-imbalance-in-shared-mobility-cost-and-consequence-72a50007831b, **2020**.

[24] R. Alvarez-Valdes, J. M. Belenguer, E. Benavent, J. D. Bermudez, F. Muñoz, E. Vercher, F. Verdejo, "Optimizing the level of service quality of a bike-sharing system", *Omega* **2016**.

[25] A. Pal, Y. Zhang, "Free-floating bike sharing: Solving real-life large-scale static rebalancing problems", *Transportation Research Part C: Emerging Technologies* **2017**.

[26] J. Zhang, M. Meng, Z. David, "A dynamic pricing scheme with negative prices in dockless bike sharing systems", *Transportation Research Part B: Methodological* **2019**.

[27] A. Mourad, J. Puchinger, C. Chu, "A survey of models and algorithms for optimizing shared mobility", *Transportation Research Part B: Methodological* **2019**.

[28] S. Shaheen, C. Bell, A. Cohen, B. Yelchuru, B. A. Hamilton, et al., Travel behavior: Shared mobility and transportation equity, tech. rep., **2017**.

[29] S. Reiss, K. Bogenberger, "A Relocation Strategy for Munich's Bike Sharing System: Combining an operator-based and a user-based Scheme", *Transportation Research Procedia* **2017**.

[30] T. Raviv, M. Tzur, I. A. Forma, "Static repositioning in a bike-sharing system: models and solution approaches", *EURO Journal on Transportation and Logistics* **2013**.

[31] I. A. Forma, T. Raviv, M. Tzur, "A 3-step math heuristic for the static repositioning problem in bike-sharing systems", *Transportation Research Part B: Methodological* **2015**.

[32] J. Schuijbroek, R. Hampshire, W.-J. van Hoeve, "Inventory rebalancing and vehicle routing in bike sharing systems", *European Journal of Operational Research* **2017**.

[33] M. Benchimol, P. Benchimol, B. Chappert, A. de la Taille, F. Laroche, F. Meunier, L. Robinet, "Balancing the stations of a self service "bike hire" system", *RAIRO - Operations Research* **2011**.

[34] D. Chemla, F. Meunier, R. Wolfler Calvo, "Bike sharing systems: Solving the static rebalancing problem", *Discrete Optimization* **2013**.

[35] M. Rainer-Harbach, P. Papazek, G. R. Raidl, B. Hu, C. Kloimüllner, "PILOT, GRASP, and VNS approaches for the static balancing of bicycle sharing systems", *Journal of Global Optimization* **2014**.

[36] G. Erdoğan, G. Laporte, R. Wolfler Calvo, "The static bicycle relocation problem with demand intervals", *European Journal of Operational Research* **2014**.

[37] M. Dell'Amico, E. Hadjicostantinou, M. Iori, S. Novellani, "The bike sharing rebalancing problem: Mathematical formulations and benchmark instances", *Omega* **2014**.

[38] M. Dell'Amico, M. Iori, S. Novellani, T. Stützle, "A destroy and repair algorithm for the Bike sharing Rebalancing Problem", *Computers Operations Research* **2016**.

[39] W. Szeto, Y. Liu, S. C. Ho, "Chemical reaction optimization for solving a static bike repositioning problem", *Transportation Research Part D: Transport and Environment* **2016**.

[40] S. C. Ho, W. Szeto, "Solving a static repositioning problem in bike-sharing systems using iterated tabu search", *Transportation Research Part E: Logistics and Transportation Review* **2014**.

[41] S. C. Ho, W. Szeto, "A hybrid large neighborhood search for the static multi-vehicle bike-repositioning problem", *Transportation Research Part B: Methodological* **2017**.

[42] Y. Wang, W. Szeto, "An enhanced artificial bee colony algorithm for the green bike repositioning problem with broken bikes", *Transportation Research Part C: Emerging Technologies* **2021**.

[43] J. Pfrommer, J. Warrington, G. Schildbach, M. Morari, "Dynamic Vehicle Redistribution and Online Price Incentives in Shared Mobility Systems", *IEEE Transactions on Intelligent Transportation Systems* **2014**.

[44] C. Contardo, C. Morency, L.-M. Rousseau, “Balancing a Dynamic Public Bike-Sharing System”, *CIRRELT:Montreal* **2012**.

[45] C. Kloimüllner, P. Papazek, B. Hu, G. R. Raidl, “Balancing bicycle sharing systems: an approach for the dynamic case”, *Springer* **2014**.

[46] S. Ghosh, P. Varakantham, Y. Adulyasak, P. Jaillet, “Dynamic Repositioning to Reduce Lost Demand in Bike Sharing Systems”, *Journal of Artificial Intelligence Research* **2017**.

[47] D. Zhang, C. Yu, J. Desai, H. Lau, S. Srivathsan, “A time-space network flow approach to dynamic repositioning in bicycle sharing systems”, *Transportation Research Part B: Methodological* **2017**.

[48] F. Chiariotti, C. Pielli, A. Zanella, M. Zorzi, “A Dynamic Approach to Rebalancing Bike-Sharing Systems”, *Sensors* **2018**.

[49] C. Shui, W. Y. Szeto, “Dynamic green bike repositioning problem – A hybrid rolling horizon artificial bee colony algorithm approach”, *Transportation Research Part D-transport and Environment* **2018**.

[50] S. Datner, T. Raviv, M. Tzur, D. Chemla, “Setting Inventory Levels in a Bike Sharing Network”, *Transportation Science* **2019**.

[51] Y. Liu, W. Szeto, S. C. Ho, “A static free-floating bike repositioning problem with multiple heterogeneous vehicles, multiple depots, and multiple visits”, *Transportation Research Part C: Emerging Technologies* **2018**.

[52] M. Du, L. Cheng, X. Li, F. Tang, “Static rebalancing optimization with considering the collection of malfunctioning bikes in free-floating bike sharing system”, *Transportation Research Part E: Logistics and Transportation Review* **2020**.

[53] G. Ma, B. Zhang, C. Shang, Q. Shen, “Rebalancing stochastic demands for bike-sharing networks with multi-scenario characteristics”, *Information Sciences* **2021**.

[54] B. Zhang, X. Li, F. Saldanha-da-Gama, “Free-floating bike-sharing systems: New repositioning rules, optimization models and solution algorithms”, *Information Sciences* **2022**.

[55] L. Caggiani, R. Camporeale, M. Ottomanelli, W. Y. Szeto, “A modeling framework for the dynamic management of free-floating bike-sharing systems”, *Transportation Research Part C: Emerging Technologies* **2018**.

[56] J. Warrington, D. Ruchti, “Two-stage stochastic approximation for dynamic rebalancing of shared mobility systems”, *Transportation Research Part C: Emerging Technologies* **2019**.

[57] X. Luo, L. Li, L. Zhao, “Dynamic Repositioning in Free-Floating Bike Sharing Systems Using Approximate Dynamic Programming”, **2021**.

[58] B. A. Neumann-Saavedra, D. C. Mattfeld, M. Hewitt, “Assessing the operational impact of tactical planning models for bike-sharing redistribution”, *Transportation Research Part A: Policy and Practice* **2021**, DOI 10.1016/j.tra.2021.06.003.

[59] M. Barth, M. Todd, “Simulation model performance analysis of a multiple station shared vehicle system”, *Transportation Research Part C: Emerging Technologies* **1999**.

[60] A. Kek, R. Cheu, M. Chor, “Relocation Simulation Model for Multiple-Station Shared-Use Vehicle Systems”, *Transportation Research Record* **2006**.

[61] A. G. Kek, R. L. Cheu, Q. Meng, C. H. Fung, “A decision support system for vehicle relocation operations in carsharing systems”, *Transportation Research Part E: Logistics and Transportation Review* **2009**.

[62] R. Nair, E. Miller-Hooks, “Fleet Management for Vehicle Sharing Operations”, *Transportation Science* **2011**.

[63] E. M. Cepolina, A. Farina, “A new shared vehicle system for urban areas”, *Transportation Research Part C: Emerging Technologies* **2012**.

[64] D. R. Jorge, G. H. A. Correia, C. Barnhart, “Comparing Optimal Relocation Operations With Simulated Relocation Policies in One-Way Carsharing Systems”, *IEEE Transactions on Intelligent Transportation Systems* **2014**.

[65] S. Weikl, K. Bogenberger, “A practice-ready relocation model for free-floating carsharing systems with electric vehicles – Mesoscopic approach and field trial results”, *Transportation Research Part C: Emerging Technologies* **2015**.

[66] Y. Deng, “Design and management of mobility-on-demand (MOD) transportation systems considering demand uncertainty and flexibility-a simulation-based approach”, **2015**.

[67] N. Jian, D. Freund, H. M. Wiberg, S. G. Henderson, “Simulation optimization for a large-scale bike-sharing system”, *2016 Winter Simulation Conference (WSC)* **2016**.

[68] K. A. Marczuk, H. S. Soh, C. M. Azevedo, D.-H. Lee, E. Frazzoli, “Simulation Framework for Rebalancing of Autonomous Mobility on Demand Systems”, *MATEC Web of Conferences* **2016**, DOI 10.1051/matecconf/20168101005.

[69] T. Zhou, C. Osorio, E. Fields, “A data-driven discrete simulation-based optimization algorithm for large-scale two-way car-sharing network design”, *Massachusetts Inst. Technol. Boston MA USA Tech. Rep* **2017**.

[70] H. R. Gómez Márquez, R. López Bracho, A. Ramirez-Nafarrate, “A simulation-optimization study of the inventory of a bike-sharing system: The case of Mexico City Ecobici’s system”, *Case Studies on Transport Policy* **2021**.

[71] Y. Jin, C. Ruiz, H. Liao, “A simulation framework for optimizing bike rebalancing and maintenance in large-scale bike-sharing systems”, *Simulation Modelling Practice and Theory* **2022**.

[72] Sid Dhuri, Spatial Data Analysis With Hexagonal Grids, https://medium.com/swlh/spatial-data-analysis-with-hexagonal-grids-961de90a220e, **2020**.

[73] VergelijkScooterverzekering, Welke deelscooters zijn er inmiddels allemaal in Nederland?, https://www.vergelijkscooterverzekering.nl/blog/welke-deelscooters-zijn-er-inmiddels-allemaal, **2021**.

[74] Daniel Corley, Visualizing Geospatial Data with Uber’s Kepler.gl, https://towardsdatascience.com/visualizing-geospatial-data-with-ubers-kepler-gl-2a437ada573d, **2020**.

[75] Eshwaran Venkat, Uber H3 for Data Analysis with Python, https://towardsdatascience.com/uber-h3-for-data-analysis-with-python-1e54acdcc908, **2021**.

7 Appendix

Several variants of the simulation model have been evaluated and compared to real data, as described in Section 4.2. The comparison of all these variations, based on the KPIs described in Section 3.2.1, is shown in the figures below. The different variants of the simulation model consist of different hexagon granularity levels being used and whether the weather type is taken into account. Lastly, the impact of basing the idle times on the day of the week or the hour of the day is considered.

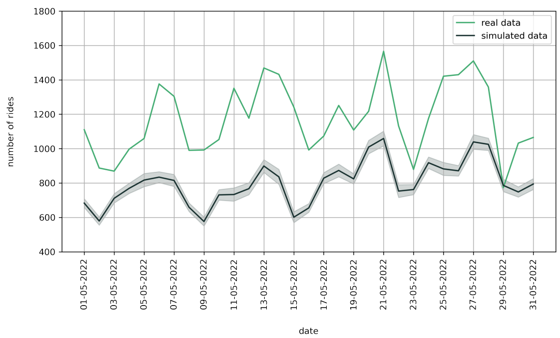

(a) Total trips

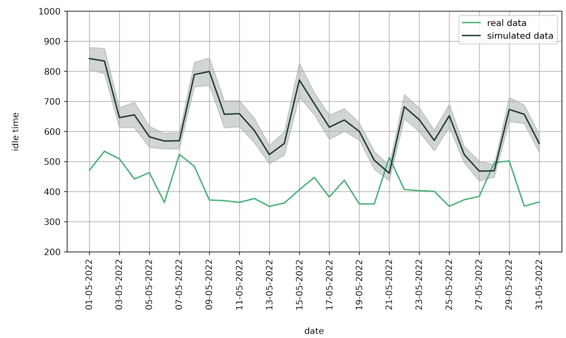

(b) Average idle time

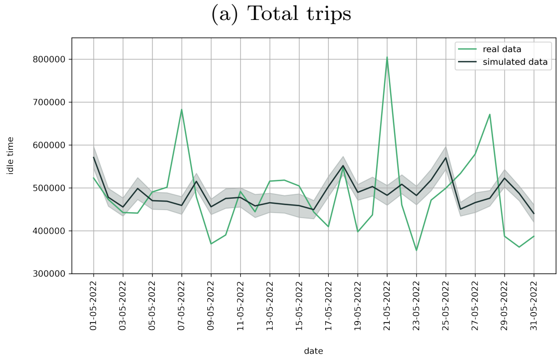

(c) Total idle time

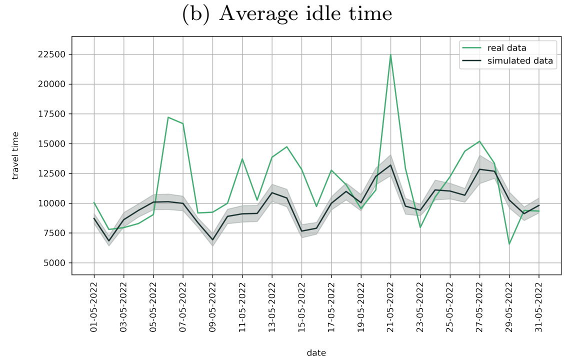

(d) Total travel time

Figure 30: Granularity level 9, type of weather excluded, idle times depend on the day of the week.

In Figure 30, the MAPE for each KPI is, respectively: 30.67%, 53.71%, 15.47%, and 19.08%. This leads to an overall MAPE of 29.73% on average.

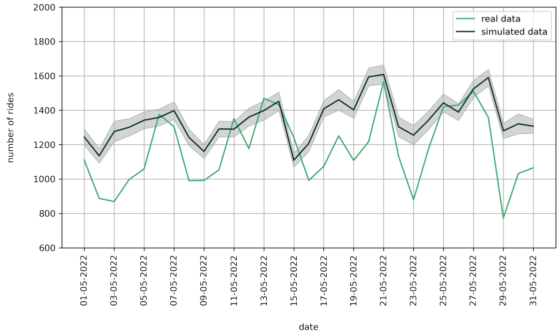

(a) Total trips

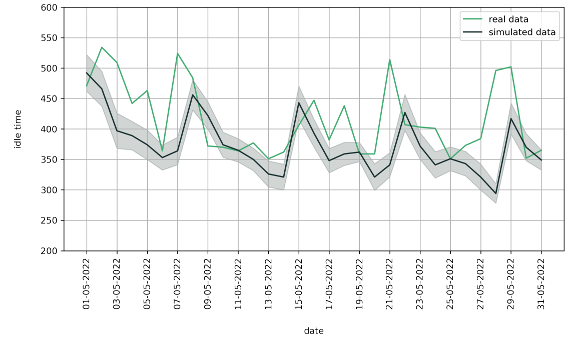

(b) Average idle time

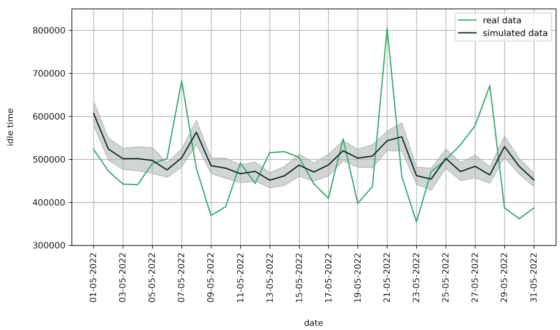

(c) Total idle time

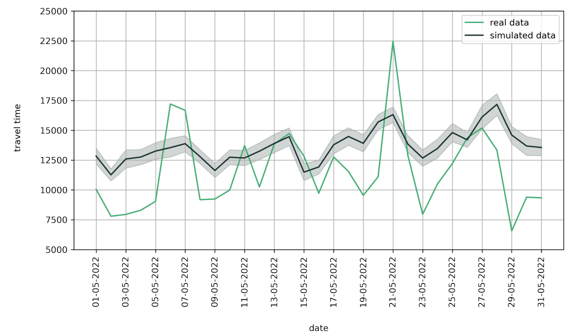

(d) Total travel time

Figure 31: Granularity level 8, type of weather excluded, idle times depend on the day of the week.

In Figure 31, the MAPE for each KPI is, respectively: 19.21%, 11.69%, 16.23%, and 30.54%. This leads to an overall MAPE of 19.42% on average.

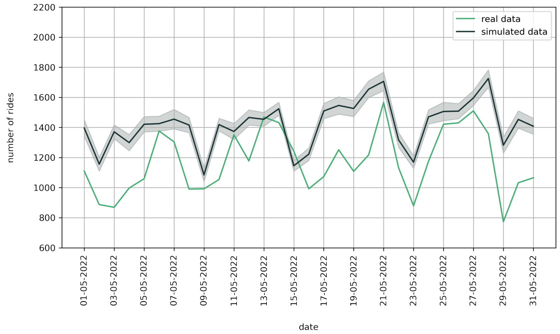

(a) Total trips

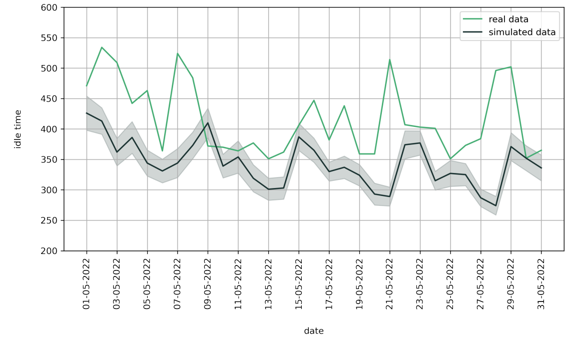

(b) Average idle time

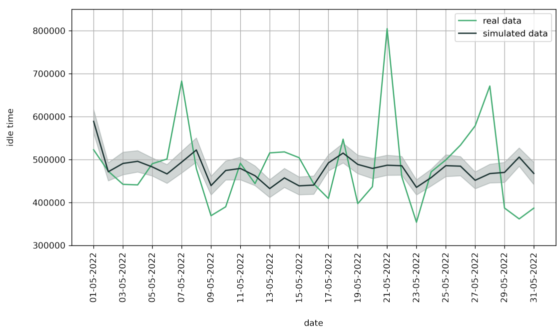

(c) Total idle time

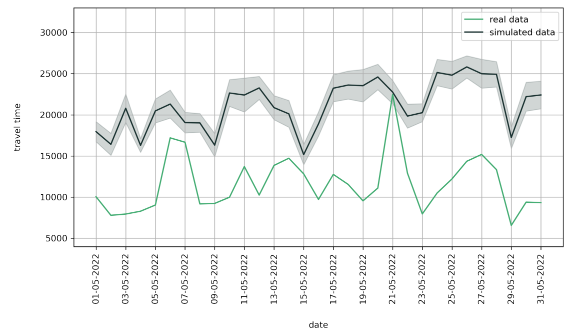

(d) Total travel time

Figure 32: Granularity level 7, type of weather excluded, idle times depend on the day of the week.

In Figure 32, the MAPE for each KPI is, respectively: 24.17%, 16.92%, 14.39%, and 93.24%. This leads to an overall MAPE of 37.18% on average.

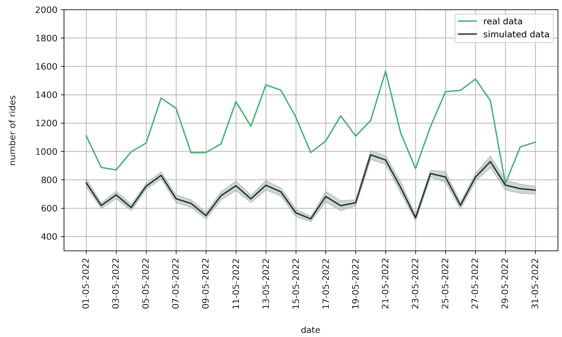

(a) Total trips

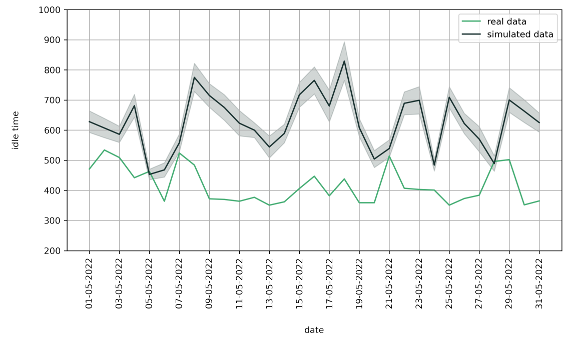

(b) Average idle time

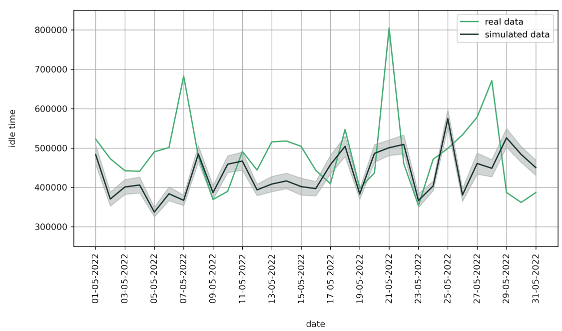

(c) Total idle time

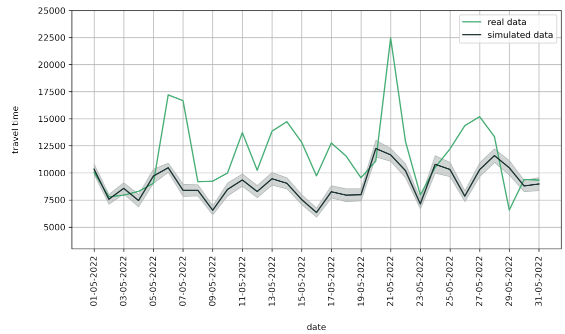

(d) Total travel time

Figure 33: Granularity level 9, type of weather included, idle times depend on the day of the week.

In Figure 33, the MAPE for each KPI is, respectively: 37.72%, 53.15%, 17.50%, and 23.37%. This leads to an overall MAPE of 32.94% on average.

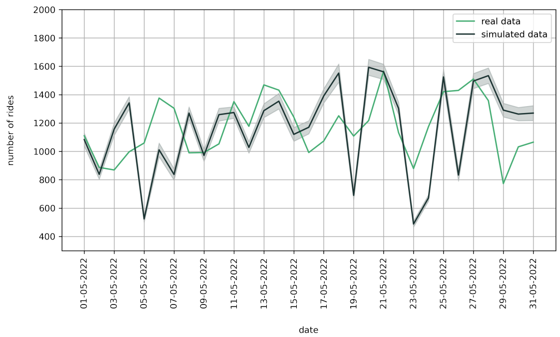

(a) Total trips

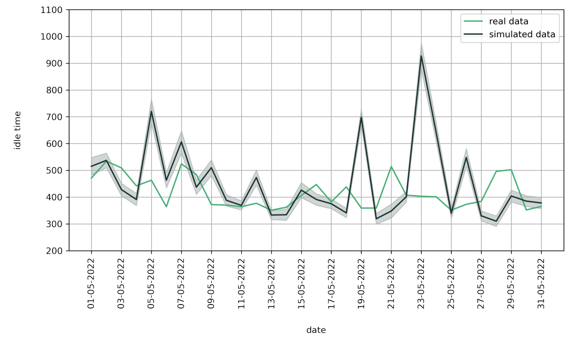

(b) Average idle time

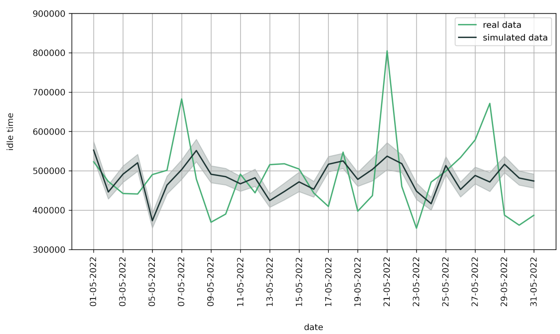

(c) Total idle time

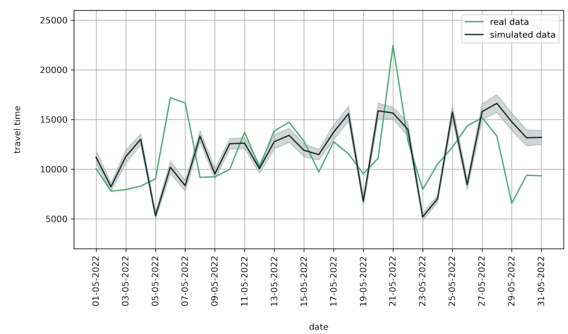

(d) Total travel time

Figure 34: Granularity level 8, type of weather included, idle times depend on the day of the week.

In Figure 34, the MAPE for each KPI is, respectively: 22.55%, 23.59%, 16.98%, and 29.03%. This leads to an overall MAPE of 23.04% on average.

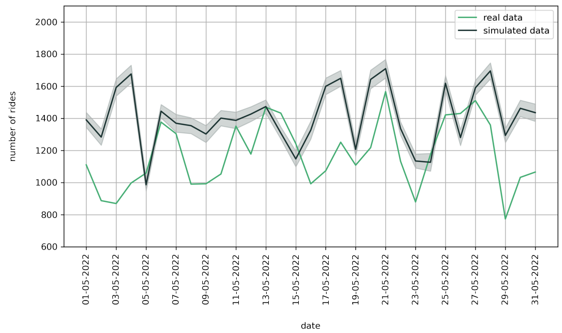

(a) Total trips

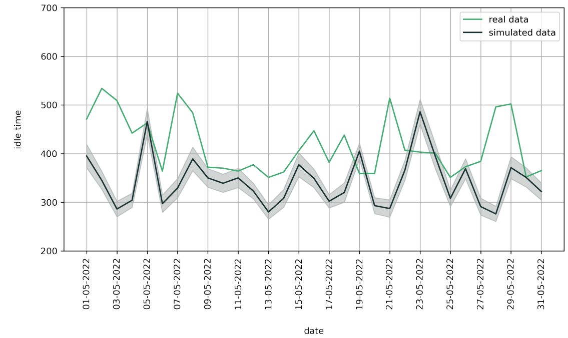

(b) Average idle time

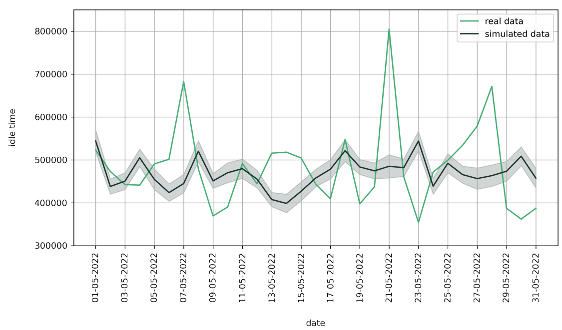

(c) Total idle time

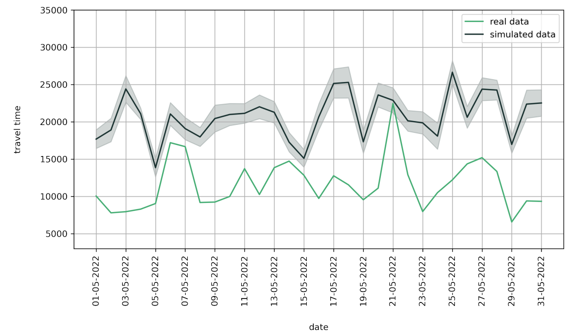

(d) Total travel time

Figure 35: Granularity level 7, type of weather included, idle times depend on the day of the week.

In Figure 35, the MAPE for each KPI is, respectively: 25.68%, 18.54%, 16.39%, and 90.27%. This leads to an overall MAPE of 37.72% on average.

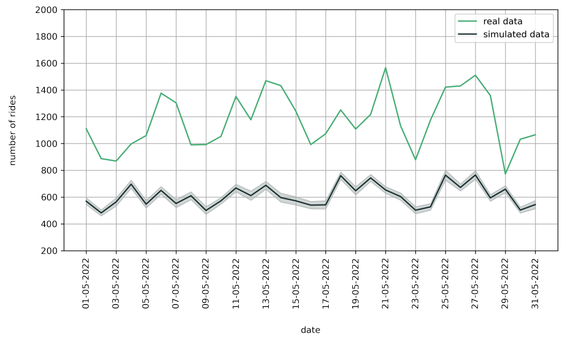

(a) Total trips

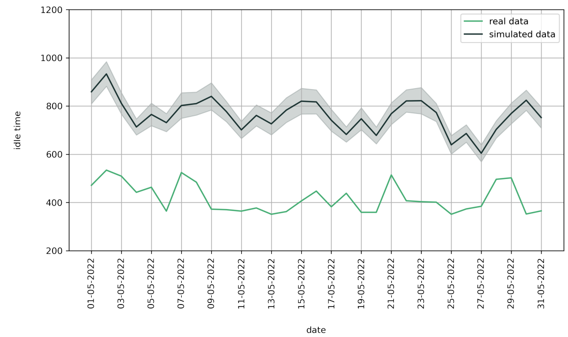

(b) Average idle time

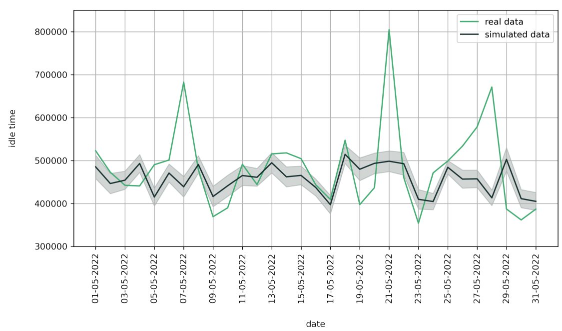

(c) Total idle time

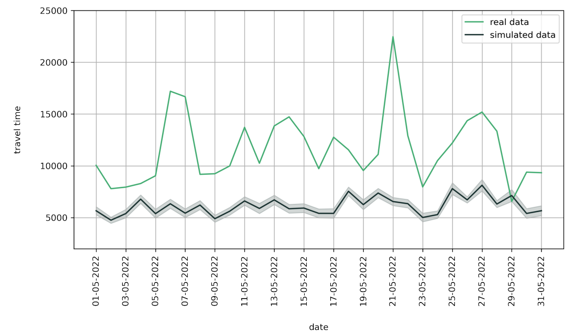

(d) Total travel time

Figure 36: Granularity level 8, type of weather excluded, idle times depend on the hour of the day.

In Figure 36, the MAPE for each KPI is, respectively: 46.90%, 85.65%, 12.53%, and 44.48%. This leads to an overall MAPE of 47.39% on average.

Research paper

Simulation-Based Optimization for Rebalancing the Fleet of Vehicles in Free-Floating Shared Mobility Systems

Stijn Mijnster[1], Andi Aliko[2], Bilge Atasoy[1]

[1] Department of Maritime and Transport Technology of Faculty Mechanical, Maritime and Materials Engineering of Delft University of Technology
[2] Felyx Netherlands B.V.
S.Mijnster@student.tudelft.nl, A.Aliko@felyx.com, B.Atasoy@tudelft.nl

A R T I C L E I N F O

Keywords:
Shared mobility
Operator-based rebalancing
Free-floating system
Discrete-event simulation
Simulation-based optimization

Abstract

In recent years, shared mobility systems have had a growing presence in cities all over the world. This is understandable given its numerous advantages such as the reduced need for personal vehicle ownership, reduced traffic congestion and emissions, increased parking efficiency, and cost savings for users. Overall, shared mobility systems offer the potential to revolutionize transportation, providing individuals with more options and helping to create more sustainable, livable cities. For shared mobility systems to fully deliver their benefits, vehicle availability must be maintained at the right place and time. If the vehicle distribution is not optimal, it may lead to overcrowding and shortages which in turn will discourage usage and lead to reduced revenues for the operator. Therefore, ensuring proper balancing of supply and demand is crucial for the success of the shared mobility service. One way to balance supply and demand is through physically rebalancing vehicles within the service area. In this study, a simulation-based optimization model is created and used to determine the optimal rebalancing operations while quantifying system improvement. A case study is conducted using real data from a moped-sharing provider to examine the impact of rebalancing operations. The results demonstrate a potential increase in profit of up to 2.06%. By performing the recommended rebalancing actions several times a week in each city where the operator is active, a significant amount of extra profit can be made. This additional profit will even rise as the usage of shared mobility rises in general.

1 Introduction

The shared use of vehicles, also known as shared mobility, has grown significantly in recent years. Shared mobility operators pop up all over the world, particularly in larger cities, providing a variety of vehicles to enable users to gain short-term access to transportation modes on an as-needed basis (Shaheen and Chan, **2016**). Examples of these sharing vehicles include cars, bicycles, mopeds, and most recently, scooters. According to some, vehicle sharing brings various advantages, including the ability for individuals to enjoy the benefits of private vehicle use without the cost and burdens of ownership (e.g., fuel, maintenance, insurance). Additionally, by sharing vehicles, fewer are required, resulting in fewer resources devoted producing them. And thirdly, shared mobility is often described as a new sustainable travel mode with low economic and environmental impact that reduces travel times on congested roads and speed up short distance trips (Bozzi and Aguilera, **2021**).

According to a McKinsey report this year, the spending on shared-mobility services, depending on customer acceptance, regulations in each country, and the progress of technology, could reach \$500 billion to \$1 trillion in 2030 [3]. Additionally, they conducted a consumer survey which reveals that 70 percent of the respondents are willing to use shared mobility vehicles for their commute [4].

From the operator's perspective, managing a shared mobility fleet involves a few interesting challenges, but there's one in particular that keeps the fleet operators occupied: maintaining a constant supply of vehicles in the right places (Trautmann and Gnägi, **2022**). Conveniently located vehicles are a defining feature of the service's user experience. Ideally, a user should always be able to find a vehicle nearby. But unfortunately for operators, the vehicles won't perfectly distribute themselves. Users often pick up vehicles at busy locations and park them in low-demand areas. This makes them out of reach for most other users, reducing their usage. If left unchecked over a period of time, the fleet as a whole can become imbalanced, with an excess of vehicles in low-demand areas and a deficiency in high-demand areas [6]. Some degree of imbalance is unavoidable. Every time a user makes a ride, they shift the spatial distribution of supply, and typically not towards the spatial distribution of demand. Operators should endeavor to minimize it, since a severe fleet imbalance can have overwhelming negative effects on ridership and revenue, eventually making it unfeasible to continue operating [7].

There are various solutions to this problem, including both hardware- and software-driven approaches. The most straightforward solution to any demand problem is to simply increase supply. Adding more vehicles to the fleet is likely to mean more will be available where and when they're needed. However, simply increasing the number of available vehicles in the fleet without optimizing the distribution, may quickly turn into loss of capital and decrease of profitability (Alvarez-Valdes et al., **2016**). Another solution is by redistributing the fleet across the city, also known as 'rebalancing', where vehicles are picked up in low-demand areas and moved to high-demand areas (Pal and Zhang, **2017**). This is often done with a van or a trailer that can transport a reasonable amount of vehicles at once. Manual relocations are fairly expensive, so they should only be carried out when unavoidable or when they'll generate a net-positive improvement on the performance of the fleet to justify the cost. A better distribution of the fleet can also be achieved by incentivizing users, eliminating the need for more vehicles or manual relocations. Such as by creating pickup and drop-off incentive zones in areas with low and high demand, respectively, or by discounting certain rides to make them more appealing (Zhang et al., **2019**).

Data-driven algorithms can be used to determine the optimal number of vehicles the fleet should have, where vehicles should be picked up and moved to, or where incentives and discounts should be applied. For the past decade, researchers have investigated this area of study extensively, trying to improve these algorithms to be faster and more accurate (Mourad et al., **2019**). To contribute to this field of research by developing an algorithm that reduces the supply and demand imbalance, one must first comprehend the current state of the fleet. Additionally, in order to properly configure the algorithm's parameters, a thorough understanding of the user's behavior is required (Shaheen et al., **2017**).

In this study, focus is on rebalancing the vehicles by physically relocating them throughout the city. To solve this vehicle rebalancing problem, past research in this field has mainly been devoted to developing analytical optimization models that determine the required rebalancing actions to shift the current distribution of the vehicles to a so-called target distribution. This target distribution would better meet the demand at that moment in time. However, the distribution of the vehicles should, ideally, not just meet the demand at that time, but also the demand over a longer period of time. As systems get more complex, estimating the behavior over a longer period of time can be difficult. Nice-form analytical models are hard to define and do not accurately capture the behavior of the system anymore and these systems may even be referred to as a 'black box'. Simulation techniques are commonly used in these situations because they are much better at taking into account the intricate interactions between supply and demand. They can be used to evaluate the system and even compare design alternatives and identify the best design among them. However, if the number of design alternatives is very large or infinite, simulation can be both expensive and time-consuming. To overcome this problem, a combination of simulation and optimization techniques can be used to determine the best design without evaluating all design alternatives. According to Zhou et al. (**2017**), combining simulation models with optimization techniques, also known as simulation-based optimization, is an innovative and promising area of future research. Simulation-based optimization involves the search for those specific settings of the input parameters such that an objective, which is a function of the simulation output, is maximized or minimized. Whereas simulation models are effective in imitating reality by taking into account uncertainties and randomness, optimization models can quickly and accurately reach optimal solutions; the advantages of both worlds are now combined. This paper shows the potential of simulation-based optimization for operator-based rebalancing in the free-floating vehicle sharing market and focuses on answering the following question:

How does the implementation of simulation and optimization techniques affect the performance of rebalancing operations in shared mobility systems?

To address this question, a discrete-event simulation is created and integrated with an optimization model to determine the optimal rebalancing operations while evaluating system improvement. The rest of the paper is structured as follows: Section 2 reviews the related literature. Section 3 describes the Vehicle Rebalancing Problem followed by Section 4 which presents the methodology for both the discrete-event simulation and the simulation-based optimization model. Section 5 depicts a case study based on real-world data from a shared moped provider, while conclusions are presented in Section 6.

2 Related literature

Barth and Todd (**1999**) developed a queuing-based discrete event simulation model that included relocations and a number of input parameters that allowed different scenarios to be evaluated. Three ways of deciding when relocations should be performed were presented: 'Static relocation' based on immediate needs in a station; 'Historical predictive relocation', which uses knowledge of expected future demand, looking 20 minutes into the future, and 'Exact predictive relocation' that can be used if perfect knowledge of future demand is available, which is impossible in the real world. The model was applied to a community in Southern California and some measures of effectiveness were calculated. The simulation model is similar to the one in this paper, but they did not develop an optimization model or ways of combining both optimization and simulation. Later, Kek et al. (**2006**, **2009**) developed an optimization model and a simulation model, but in their work only the optimization models allow for determining the relocation operations. The simulation model is just used to evaluate the performance of the systems when the relocation operations determined by the optimization model are performed. Nair and Miller-Hooks (**2011**) continued exploring optimization methods and proposed a stochastic mixed-integer programming (MIP) model to optimize vehicle relocations, which has the advantage of considering demand uncertainty. However, they did not develop a simulation model.

Cepolina and Farina (**2012**) propose a methodology, based on the Simulated Annealing (SA) algorithm to optimize the fleet distribution of a station-based car-sharing system. The reason for this is that there is no analytical expression for the cost function, so the chances are high that a local optimum is reached instead of a global optimum, and the search space is extremely large. The methodology includes a simulation model of the proposed transport system which allows one to track the second-by-second activity of each user, as well as the second-by-second activity of each vehicle. The cost function consisting of the transport management cost (i.e. the cost of vehicles) and the cost to the customer (i.e. the total customer waiting time) is minimized by explicitly simulating the arrival of the users, the departure of the vehicles from the stations and the arrival of the vehicles at the stations. Jorge et al. (**2014**) present two independent tools that can be combined: a mathematical model for optimal vehicle relocation, and a discrete-event time-driven simulation model with several real-time relocation policies integrated. Results show that relocating vehicles, using any of the methods developed, can produce significant increases in profit. Well, the developed simulation model here is only used for evaluating the rebalancing policies.

Weikl and Bogenberger (**2015**) provided relocation strategies for free-floating systems for pick up and drop-off. They combine a macroscopic relocation optimization policy of moving vehicles between zones, with a rule based heuristic for station to station relocations. They make use of a historical data analysis that generates the input for the calculation of a target vehicle distribution for different target periods. If vehicle supply and demand deviate from each other, an optimization model is used to calculate profit maximizing zone to zone relocations. The relocation strategies have been tested in a real-world setting rather than in a simulation. Deng (**2015**) developed a decision support tool to assist with determining the optimal fleet configuration of a MoD system accounting for stochastic demand and the effect of conducting vehicle distribution as part of daily operations. An optimization problem is defined to find the optimal fleet configuration is terms of minimizing cost and satisfying a certain level of service. A discrete-event simulator (DES) that includes a sub-optimization model to calculate hourly rebalancing schemes is built to estimate the performance of a given configuration. Finally, an algorithm is devised that combines Particle Swarm Optimization (PSO) and Optimal Computation Budget Allocation (OCBA) techniques to efficiently search the design and decision space.

Jian et al. (**2016**) use DES to model a station-based bike-sharing system. They tackle the rebalancing problem over bikes and docks as a simulation-optimization problem. Ideally, they would apply standard simulation-optimization methods, such as stochastic gradient-search and random search, to solve the problem, but this seems computationally infeasible. Instead, they develop heuristic search procedures that use statistics from a single simulation run in order to update the allocation of bikes and docks between stations. In each iteration they generate a trial solution and evaluate it with the DES model. If the trial solution improves the objective, then they move to that solution, otherwise they stay at the last solution. They do not claim that they find local or global optima, but instead see the value of these algorithms in the improvements they make in performance relative to that of starting solutions. Marczuk et al. (**2016**) develop several optimization models for three rebalancing policies within car-sharing systems: i) no rebalancing (baseline), ii) offline rebalancing, and iii) online rebalancing. The performance of the three policies are then evaluated using the simulation program SimMobility. Zhou et al. (**2017**) propose a car-sharing optimization problem also as a simulation-optimization (SO) problem. Here, no analytical expression of the objective function is available, hence traditional (analytical) discrete optimization algorithms cannot be used. A novel metamodel is formulated, which is based on a MIP formulation. The metamodel is embedded within a general-purpose discrete SO algorithm. The combination of the problem-specific analytical MIP with a general-purpose SO algorithm enables to address high-dimensional problems and become computationally efficient. More generally, the information provided by the MIP to the SO algorithm enables it to exploit problem-specific structural information. Hence, the simulator is no longer treated as a black box.

Gómez Márquez et al. (**2021**) develop a simulation-optimization framework to determine the bike inventory for stations in a large-scale bike-sharing system. The framework helps to optimize both the bike inventory at the beginning of the day, which is the focus of static rebalancing, and the bike inventory throughout the day, which is the focus of dynamic rebalancing. They implement several simulation-optimization methods including nested partitions (NP), interactive particle algorithm (IPA), cross entropy, and discrete simultaneous perturbation stochastic approximation (DSPSA) and find that IPA provides good solutions within reasonable computing time. Jin et al. (**2022**) propose a simulation framework for evaluating different rebalancing and maintenance strategies to model the daily operations of large-scale bike-sharing systems with docking stations. The framework can be integrated with any multi-vehicle static or dynamic rebalancing optimization model. An optimization model solved by an enhanced k-means clustering method (EKM) and an Ant Colony Optimization (ACO) algorithm is provided as an example for demonstrating such integration. Although the proposed simulation framework is developed for bike-sharing systems, it can be easily modified for modeling other transportation systems with non-floating stations (e.g. electrical bikes and scooters).

In most of the mentioned studies, simulation and optimization techniques are used to determine and evaluate rebalancing operations within various shared mobility systems. In some cases, simulation is just used to evaluate particular rebalancing strategies, while in other cases, simulation and optimization models are really integrated. Additionally, in some of these studies, historical data of trips is used to determine a target distribution, which is then compared to the actual distribution of the vehicles. Rebalancing actions are then suggested to reduce this imbalance. In this study, historical data of trips is also used, but not directly to set up a target distribution. First, the historical data is used as an input for a simulation model to estimate how the system is likely to behave over a given period of time. Secondly, the occurrences in the simulation are then used in an optimization model to determine the optimal rebalancing actions. The distinction is that in this case, not only is historical data examined, but it is also used to predict future events based on the current distribution of the vehicles. The simulation and optimization models are thereby integrated. In the majority of the mentioned studies where a simulation model is used, demand prediction is employed. Predicting the demand is challenging, and a lot of information is typically lacking here. This is because unmet demand, also known as latent or censored demand, is not taken into consideration. The demand prediction is primarily based on trips that took place because there were vehicles available. Therefore, in this study it is decided to only use existing data, such as the trips and idle times that actually took place. To be more precise, in most simulation models the vehicles stand idle until demand pops up nearby a vehicle to make a ride. The vehicles in the simulation model in this study stand idle for a certain idle time, sampled from the idle time distribution associated with its location before it will make a new ride. As a result,

Table 1: Summary of simulation-based optimization rebalancing problem literature.

Reference	Simulation model	Optimization model	Methodology	Vehicle	Type
Barth and Todd (1999)	✓	✗	evaluation by simulation	CS	station-based
Kek et al. (2006, 2009)	✓	✓	evaluation by simulation	CS	station-based
Nair and Miller-Hooks (2011)	✗	✓	-	CS	station-based
Jorge et al. (2014)	✓	✓	evaluation by simulation	CS	station-based
Weikl and Bogenberger (2015)	✗	✓	-	CS	free-floating
Cepolina and Farina (2012)	✓	✓	simulation apart from optimization	CS	station-based
Deng (2015)	✓	✓	simulation and optimization integrated	CS	station-based
Jian et al. (2016)	✓	✓	simulation and optimization integrated	BS	station-based
Marczuk et al. (2016)	✓	✓	evaluation by simulation	CS	station-based
Zhou et al. (2017)	✓	✓	simulation and optimization integrated	CS	station-based
Gómez Márquez et al. (2021)	✓	✓	simulation and optimization integrated	BS	station-based
Jin et al. (2022)	✓	✓	simulation and optimization integrated	BS	station-based
This paper	✓	✓	simulation and optimization integrated	MS	free-floating

Explanation of terms: CS = car-sharing, BS = bike-sharing, MS = moped-sharing

the proposed approach enables to replicate the real system reasonably well without relying on predictions of demand. Table 1 summarizes the literature together with the contribution of this work.

Finally, the proposed simulation-based optimization model in this study is tested on real data and is applicable to various free-floating shared mobility systems as well as to any city where the operator have been operational for a sufficient length of time that enough data is collected, with one year being preferred.

3 Problem description

We consider a vehicle rebalancing problem in a free-floating shared mobility system from the perspective of the operator.

We assume a shared mobility operator providing shared *vehicles*; vehicles that have been modified in such a way that they can be accessed and used by anyone via an app. Examples of these vehicles are bicycles, cargo bikes, mopeds, and cars. Individuals who use these vehicles are referred to as *users*. The vehicles are provided in cities where the shared mobility operator is active. These cities have a *service area*; a GPS-based virtually confined area where the vehicles can be used. The user can make a *ride*; a ride can start and end at any location within the service area. The number of active vehicles and their related locations, we assume as the *supply*. The number of individuals who want to take a ride and their locations are referred to as the *demand*. To match the locations of the vehicles with the locations of the potential users, in other words, to match supply and demand, *rebalancing vehicles* are used. These rebalancing vehicles pick up the shared vehicles and relocate them within the service area. Vans or trailers that can transport a reasonable amount of vehicles at once are commonly used as rebalancing vehicles to perform these *rebalancing actions*.

After a user completes their ride, it leaves the vehicle at its destination. The vehicle will then be available to other users to make a ride. The elapsed time between the end of the ride and the start of a new ride for the same vehicle is referred to as the *idle time*. The idle time is linked to the end location of the first ride. The pickup location of the vehicle by the user depends partially on where the user is situated, but of course also on the supply of vehicles at that time of the day. The user's intent determines the drop-off location after a ride. The movement of the vehicle from the pickup location to the drop-off location is assumed as the *transition*. This transition has a certain duration, which we refer to as the *travel time*. This travel time depends on the route taken by the user and may vary due to traffic at that time of the day.

An example of a shared mobility system where rebalancing actions are performed is illustrated in Figure 1. The black border line denotes the service area. The grey circles depict users that want to make a ride. The green squares represent vehicles that are making a ride within the service area. These rides are indicated by a thin dotted line with a green location symbol as the start point and a red location symbol as the end point. The orange squares represent vehicles that are currently not in use and are therefore standing idle. The blue rectangle represents a rebalancing vehicle that is performing a rebalancing action, which is indicated with a thick dotted line with again a green location symbol as the start point and a red location symbol as the end point.

The goal is to reduce the imbalance between vehicle supply and demand by using rebalancing vehicles that perform rebalancing actions.

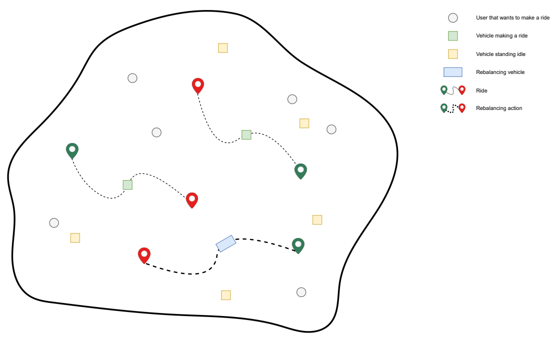

Figure 1: An illustration of a shared mobility system where rebalancing actions are performed.

4 Proposed methodology

In this section, the simulation approach is described, followed by a detailed explanation of the simulation-based optimization model.

4.1 Simulation

A discrete-event simulation (DES) model is created to simulate the behavior of vehicles in a shared mobility system over time. The simulation model is written with Python 3.18.13 along with the Salabim 22.0.7 Python package. In short, the simulation model is constructed as follows. There is one class, which is the vehicle, that follows a certain process. This process consists of an initial step followed by three sequential events that are repeated till the simulation runtime is over. During the simulation, all information about these events is monitored, including the timestamps, locations, and vehicles involved, which is later used in the integration with an optimization model. Based on the user's input about the city to be simulated and the start moment of the simulation, disaggregated historical data is retrieved from actual rides. From this data, the model inputs are determined and are as follows.

<u>Idle times</u> | The elapsed time between the end of a ride and the start of a new ride for the same vehicle is referred to as the idle time. This idle time is linked to the end time and location of the first ride.

<u>Transitions</u> | Where rides go to, based on its start location, is referred to as the transition. This transition is linked to the start location of the ride.

<u>Travel times</u> | The duration of rides, based on its start and end location, is referred to as the travel time. This travel time is linked to the route between the start and end location of the ride.

The historical data usually includes the exact location of the start and end points of rides, which are retrieved by GPS sensors located on the vehicles. The sensors' output is generally in geographic coordinates (latitude/longitude). Analyzing this data based on its exact location is both difficult and expensive. As a result, it is common practice to enclose this data in grid cells. These grid cells aggregate the underlying data points, which are then represented by a 'small' area. In this way, analyses can be carried out much easier and more efficient. Shared mobility systems rely on accurate mapping of geographical areas for their services. Therefore, it is crucial to use a grid map that minimizes distortion and quantization error introduced when users move through a city, which is the case with hexagonal grid cells. Uber also analyzes spatial data using hexagonal grid cells and has open-sourced it's hexagon mapping library H3, which can be used for this [26].

Simulation process
The initial step is to retrieve the current location of all vehicles at the start moment of the simulation. Following that, at time step zero, all vehicles are created in the simulation model and given their corresponding location and vehicle ID. As a result, the vehicle distribution is exactly the same as in reality at that point in time. The second step in the process is that all vehicles are given a certain idle time. This idle time is location- and time-dependent and is drawn from a distribution. This distribution is based on data associated with the vehicle's location and the current time step in the simulation. As a result, the idle time for each vehicle will be unique. All vehicles will wait until the idle time is over before proceeding to the next event. Because each vehicle has been assigned a unique idle time, the next event for each vehicle will occur at different time steps in the simulation. For now, we will focus on a single vehicle to describe the next steps in the process of the simulation. The third step in the process, once the idle time is over, is for the vehicle to make a ride. The start location of the ride, which is the current location of the vehicle, is of course known, but the end location must be determined. The simulation model derives the end location based on data from rides with the same start location as where the vehicle is currently located and around the current time step in the simulation. It selects the end location relying on a probability distribution. Now that the end location has been determined, the vehicle will travel from its current location to the end location of the ride. Of course, this movement takes time, which must be accounted for in the simulation. Therefore, the fourth step of the process is to establish the travel time belonging to this ride. The travel time is drawn from a distribution based on data from rides with the same start and end location. In contrast to the idle time and the end location of the ride, it is assumed here that the travel time is independent of time. After this travel time, the vehicle will be at its new location, and steps two to four will be repeated until the simulation runtime is over. Although only one vehicle is considered here, all vehicles follow the same process simultaneously. A sensitivity analysis is performed, after it was decided to run the simulation model for 50 replications to obtain accurate outcomes in a reasonable amount of time. Following that, the sample mean and standard deviations are calculated.

4.2 Simulation-based optimization

As described above, based on the location of the vehicles at the start of the simulation, the shared mobility system can be simulated over a certain period of time. To better match supply and demand, rebalancing actions can be performed. These actions alter the start location of the vehicles that are rebalanced. In other words, the distribution of the fleet of vehicles throughout the service area is adjusted. In most cases, randomly moving vehicles does not improve the performance of the system. Therefore, an optimization model is used to determine the optimal rebalancing actions to better match supply and demand.

As input for the optimization model, data from the simulation is used. This data includes all rides taken during each simulation replication, as well as all idle times at each location after these simulated rides. The information on all rides is then used to compute the average number of outgoing and incoming rides per location per hour. With this, together with knowing the number of vehicles at each location at the start of the simulation, the evolution of the number of vehicles per location per hour can be determined. It is then possible to identify which locations have the greatest surplus of vehicles during the simulation period. Please take note that the number of outgoing rides per location per hour is limited by the number of vehicles available at this location at that moment in time. As a result, the identified surplus will never be negative for any location.

In the optimization model, the idle times of the vehicles per location are taken into account, as well as the magnitude of the surplus of vehicles at this location during the simulation. Vehicles are rebalanced from locations with high idle times and a large surplus of vehicles to locations with low idle times and a potential deficit of vehicles. A potential deficit is said to exist if the surplus of vehicles during the simulation approaches or even reaches zero.

The vehicle rebalancing problem is defined on a directed graph $G = (H, A)$, where the set H contains the hexagons and the set A contains the arcs. The set of arcs $A = H \times H$ consists of all feasible arcs as $A = \{(i, j) \mid i \in H, j \in H, i \neq j\}$. $H1 \subset H$ and consists of all hexagons with a surplus of vehicles greater than zero. $H2 \subset H$ and contains the hexagons with the 25% lowest average idle times, where the average idle time (determined from all simulation replications) is based on more than 100 data points. The rebalancing decision variable x_{ij} is equal to 1 if a vehicle is rebalanced from hexagon i to hexagon j and 0 otherwise. Table 2 summarises the sets, parameters, and decision variables used in the binary integer programming (BIP) formulation.

Table 2: The sets, parameters, and decision variables for the Vehicle Rebalancing Problem.

Sets	
H	Hexagons, indexed by i, j and $\|H\| = h$
A	Arcs
$H1$	Hexagons with a surplus of vehicles greater than zero
$H2$	Hexagons with the 25% lowest average idle times and based on more than 100 data points
Parameters	
$idle\_time$	Average idle time of vehicles in a hexagon
$surplus$	Surplus of vehicles in a hexagon
rc	Rebalancing capacity
Decision variables	
x_{ij}	Binary variable

The model is formulated as follows.

Minimize:

$$\sum_{i \in H1} \sum_{j \in H2} \frac{idle\_time_j * surplus_j}{idle\_time_i * surplus_i} * x_{ij} \tag{1}$$

Subject to:

$$\sum_{i \in H1} \sum_{j \in H2, j \neq i} x_{ij} = 1 \tag{2}$$

$$x_{ij} \in \{0, 1\} \quad \forall\, i, j \in A \tag{3}$$

The objective function (1) minimizes the idle time multiplied by the magnitude of the surplus of the pickup location, divided by the idle time multiplied by the magnitude of the surplus of the drop-off location. Again, the decision variable represents the optimal rebalancing action. Constraint (2) ensures that the decision variable only contains one rebalancing action. The set $H1$ ensures that vehicles can only be picked up at hexagons with vehicles available. The set $H2$ ensures that vehicles are dropped off at hexagons with the 25% lowest idle times and with an average idle time based on a sufficient number of data points. One thing to note is that if the surplus of vehicles for a particular hexagon is zero, this value is adjusted to 0.0001. This ensures that the objective function will never be zero and that therefore always a solution can be found. Constraint (3) defines x as a binary variable.

The optimal rebalancing action determined in the model causes the number of vehicles in a certain hexagon to decrease by one and to increase by one in another hexagon. The initial distribution of the vehicles over the hexagons therefore changes. This alters the calculation for determining the surplus of vehicles in the hexagons as well. By applying the change in the vehicle distribution and recalculating the surplus of vehicles in the hexagons, it is possible to run the optimization model again and determine the next optimal rebalancing action. This process can be repeated until a certain rebalancing capacity is reached. The idle time per hexagon is not dependent on the distribution of the vehicles and therefore does not change during this recalculation of the surplus of vehicles. Algorithm 1 describes this approach using pseudocode.

Algorithm 1 Determine multiple rebalancing actions

$fleet\_distribution$ = **initial fleet distribution**
$idle\_times$ = **average idle times**
$i = 0$

while $rc < i$ **do**
  $surpluses$ = **calculate surpluses**($fleet\_distribution$)
  $rebalancing\_action$ = **optimize**($surpluses, idle\_times$)
  $fleet\_distribution$ = **perform rebalancing action**($rebalancing\_action, fleet\_distribution$)
  i += 1
end while

5 Case study

A case study is conducted using real-life data from a moped-sharing provider. The dataset includes data from all trips that took place in a single city between May '21 and May '22 and contains timestamps, vehicle ids, start and end locations, and trip durations. First of all, this dataset is used to optimize the performance of the simulation model such that the results in terms of total trips, average idle time, total idle time, total travel time, and the locations visited, correspond as closely as possible to reality. The simulation model is then integrated with the optimization model in order to identify the best actions to take to potentially improve the system. The whole month of May '22 has been subjected to the simulation-based optimization model. Figure 3 shows the optimal rebalancing actions for various rebalancing capacities on the first of May.

The hexagons indicate the locations that are used throughout the simulation replications, whereas the numbers represent the initial distribution of the vehicles. The vehicles must be picked up from the hexagons with yellow marks and dropped off at the hexagons with orange marks. The opacity of the blue line reflects the number of vehicles that must be rebalanced. The higher the opacity, the more vehicles to rebalance. To view this more clearly, the rebalancing actions will, of course, also be listed in tables.

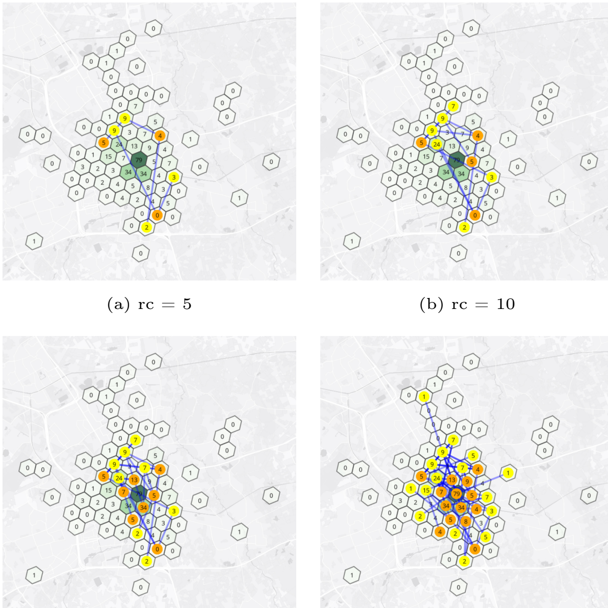

(a) rc = 5
(b) rc = 10
(c) rc = 20
(d) rc = 50

Figure 3: Rebalancing actions on May 1 with various rebalancing capacities (rc).

With the optimal rebalancing actions to perform at the start of the simulation known, a comparison can be made between the system with the initial distribution of the vehicles and the system with the distribution of the vehicles after the rebalancing actions, using the simulation model. The results for the situation on May 1 indicate that for all rebalancing capacities, the total number of trips, the total travel time, and therefore the revenue made on this day, will increase by performing the recommended rebalancing actions.

The simulation-based optimization model is used throughout the whole month of May to determine whether and how much the system may potentially improve each day. For instance, there are days when the distribution of the vehicles is more concentrated in the city center, making it difficult to increase revenue, and days when the vehicles are widely dispersed, including the city's outskirts, making it easier to generate more revenue (see Figure 5).

(a) May 17
(b) May 29

Figure 5: Rebalancing actions on two days that show contrasting behavior with rc = 20.

It is important to note that another study by this moped-sharing provider revealed that rebalancing actions affect the generated revenue not just on the day they are performed, but also on the days following. Vehicles that may have been idle for several days at their initial locations are now rebalanced to locations with higher demand. From these new locations, users take the vehicles and drive to other areas, where demand is again likely to be higher than at its initial location, resulting in several extra rides over multiple days. According to this study, the effect can be noticed for up to three days. Therefore, the improvement of the system should not be studied per day, but as a continuum over the course of three days. This implies that the distribution of the vehicles at the end of the day will serve as the starting point of the simulation on the following day (it is assumed that there's no activity during the night because only 3.48% of the rides in this city occur during these hours).

Figure 7 depicts the revenue increase in percentages over three consecutive days, where the rebalancing actions are only performed on the first day, which are in this case May 13 and May 27. The error bars indicate the standard error (SE) which is calculated from five model replications. It can be seen that there is an increase in revenue not only on the first day but also on the two following days. For the 13th of May, the total revenue increase over three days is on average 1.25%, 2.05%, 3.22%, and 4.35% for rebalancing 5, 10, 20, or 50 vehicles, respectively. For the 27th of May, this is 1.59%, 2.36%, 3.77%, and 6.57%.

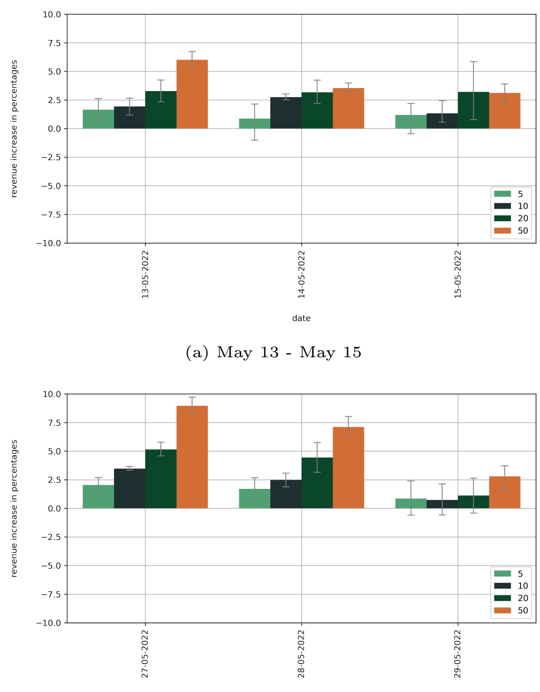

(a) May 13 - May 15

(b) May 27 - May 29

Figure 7: Revenue increase over three consecutive days due to rebalancing actions on the first day.

In order to asses whether the extra revenue generated by the rebalancing actions also results in extra profit, several costs must be considered. Based on conversations with multiple employees of the moped-sharing provider, the process of performing the rebalancing actions and the related costs are drawn out. Figure 9 depicts the final extra profit for the same three consecutive days as shown before. A comparison of the revenue and profit graphs reveals that on the first day, the profit is significantly lower, if not negative, than the revenue. This is due to the costs of the rebalancing operations, which only take place on the first day. Additionally, a certain percentage of the revenue is withheld in all cases to cover some general operational costs. For the 13th of May, the total profit increase over three days is on average 0.58%, 0.84%, 1.07%, and 0.10% for rebalancing 5, 10, 20, or 50 vehicles, respectively. For the 27th of May, this is 0.85%, 1.14%, 1.62%, and 2.06%.

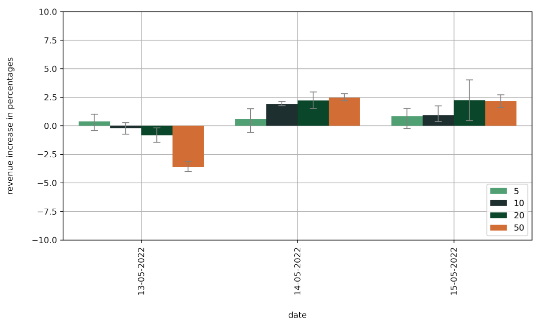

(a) May 13 - May 15

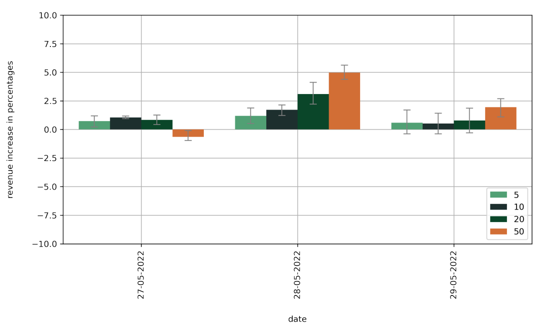

(b) May 27 - May 29

Figure 9: Revenue increase over three consecutive days due to rebalancing actions on the first day.

Rebalancing usually takes place on a regular basis, suppose around two to three times a week. The profit growth in euros can therefore rise significantly over the course of a year. Additionally, rebalancing actions can be performed in each city an operator is active in, raising this number even higher.

6 Conclusions and future research

A discrete event simulation model is created that imitates the behavior of a free-floating shared mobility system by using a hexagonal grid system. The simulator uses historical data up to the simulation's starting point as input and can therefore simulate periods from the past as well as periods from present time into the future. It runs for a sufficient number of replications and calculates the sample means and standard deviations, which are compared to real data. Following that, the simulation results are used in an optimization model. The average idle times together with the surpluses or potential deficits of vehicles throughout the simulation are determined per location. The optimal rebalancing actions then indicate that the vehicles must be rebalanced from locations with a high average idle time and a surplus of vehicles to locations with a low average idle time and a potential deficit of vehicles.

With the optimal rebalancing actions established, the results of the simulation of the system with the initial distribution of the vehicles can be compared with the results of the simulation of the system with the distribution of the vehicles after the rebalancing actions have been performed. Crucial here is that rebalancing actions affect the generated revenue not just on the day they are performed, but also on the days following. Therefore, the total extra revenue that is generated over three days, where rebalancing actions are only performed on the first day, is examined. The results from two distinct days demonstrate an increase in revenue over three days of 4.35% and 6.57% by rebalancing 50 vehicles. In order to calculate the extra profit, the costs of the rebalancing operations must, of course, be taken into account. The final results show a possible increase in profit of 1.07% and 2.06% for the two days that are examined. Rebalancing actions usually take place multiple times per week in every city an operator is active in and can therefore lead to significant profit growth in euros over the course of a year.

The results obtained with the simulation-based optimization model will vary significantly depending on the shared mobility system under consideration, as well as on the city, the season, and even the day. It is therefore challenging to assess the model's real performance. To be more explicit, if the use of shared vehicles in a shared mobility system is higher in a certain period, performing rebalancing actions will result in extra revenue than if the vehicles are used less frequently. The costs of rebalancing operations, on the other hand, will remain constant, resulting in more profit. If the use of the vehicles is low, it may be required to perform rebalancing operations to better match the demand (or increase the service level), which may result in a loss at the moment, but will eventually lead to a profitable operation if the use of the vehicles rises.

Future research should include calculating the required number of rebalancing actions to optimize system improvement, rather than doing so per 5, 10, 20, or 50 vehicles as is done right now. The model can be adjusted in such a way that it stops rebalancing as soon as the improvement yields less than a certain threshold. Another approach could be to examine the surpluses and potential deficits in the system, which are currently already included in the optimization model, and that the model stops rebalancing once this imbalance is reduced to a certain value. Furthermore, it is recommended to perform the rebalancing actions from the simulation-based optimization model using a vehicle routing algorithm. To be specific, performing the rebalancing actions can be seen as a capacitated vehicle routing problem with pickup and delivery (CVRPPD). Using this algorithm, the optimal route can be determined to perform the rebalancing actions, taking into account the capacity of the rebalancing vehicle. First of all, this algorithm is also able to better draw out the actual costs, and secondly, various rebalancing strategies may be assessed. Consider, for example, the use of different vehicles with varying capacities. Finally, combining operator-based rebalancing with user-based rebalancing could be an interesting subject for future research. User-based rebalancing can sometimes be more effective due to its low costs. However, this alone is insufficient, which is why a combination with operator-based rebalancing is required. The development of a single algorithm in which both types of rebalancing are combined might result in an even greater improvement in service quality as well as in generated profit, which is worth investigating.

References

[1] S. Shaheen, N. Chan, "Mobility and the Sharing Economy: Potential to Overcome First- and Last-Mile Public Transit Connections", *Research Papers in Economics* **2016**.

[2] A. D. Bozzi, A. Aguilera, "Shared E-scooters: a review of uses, health and environmental impacts, and policy implications of a new micro-mobility service", *Sustainability* **2021**.

[3] K. Heineke, B. Kloss, A. Mertens-von Rüden, T. Möller, C. Wiemuth, Shared mobility: Sustainable cities, shared destinies, https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/shared-mobility-sustainable-cities-shared-destinies, **2023**.

[4] K. Heineke, B. Kloss, F. Rupalla, D. Scurtu, Why micromobility is here to stay, https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/why-micromobility-is-here-to-stay, **2021**.

[5] N. Trautmann, M. Gnägi, *Operations Research Proceedings 2021*, Springer Publishing, **2022**.

[6] D. Cataneo, The smart way to balance supply and demand in a shared micromobility, https://www.electricfeel.com/blog/supply-demand-shared-micromobility, **2022**.

[7] E. Fields, Supply-demand imbalance in shared mobility: cost and consequence, https://medium.com/zoba-blog/supply-demand-imbalance-in-shared-mobility-cost-and-consequence-72a50007831b, **2020**.

[8] R. Alvarez-Valdes, J. M. Belenguer, E. Benavent, J. D. Bermudez, F. Muñoz, E. Vercher, F. Verdejo, "Optimizing the level of service quality of a bike-sharing system", *Omega* **2016**.

[9] A. Pal, Y. Zhang, "Free-floating bike sharing: Solving real-life large-scale static rebalancing problems", *Transportation Research Part C: Emerging Technologies* **2017**.

[10] J. Zhang, M. Meng, Z. David, "A dynamic pricing scheme with negative prices in dockless bike sharing systems", *Transportation Research Part B: Methodological* **2019**.

[11] A. Mourad, J. Puchinger, C. Chu, "A survey of models and algorithms for optimizing shared mobility", *Transportation Research Part B: Methodological* **2019**.

[12] S. Shaheen, C. Bell, A. Cohen, B. Yelchuru, B. A. Hamilton, et al., Travel behavior: Shared mobility and transportation equity, tech. rep., **2017**.

[13] T. Zhou, C. Osorio, E. Fields, "A data-driven discrete simulation-based optimization algorithm for large-scale two-way car-sharing network design", *Massachusetts Inst. Technol. Boston MA USA Tech. Rep* **2017**.

[14] M. Barth, M. Todd, "Simulation model performance analysis of a multiple station shared vehicle system", *Transportation Research Part C: Emerging Technologies* **1999**.

[15] A. Kek, R. Cheu, M. Chor, "Relocation Simulation Model for Multiple-Station Shared-Use Vehicle Systems", *Transportation Research Record* **2006**.

[16] A. G. Kek, R. L. Cheu, Q. Meng, C. H. Fung, "A decision support system for vehicle relocation operations in carsharing systems", *Transportation Research Part E: Logistics and Transportation Review* **2009**.

[17] R. Nair, E. Miller-Hooks, "Fleet Management for Vehicle Sharing Operations", *Transportation Science* **2011**.

[18] E. M. Cepolina, A. Farina, "A new shared vehicle system for urban areas", *Transportation Research Part C: Emerging Technologies* **2012**.

[19] D. R. Jorge, G. H. A. Correia, C. Barnhart, "Comparing Optimal Relocation Operations With Simulated Relocation Policies in One-Way Carsharing Systems", *IEEE Transactions on Intelligent Transportation Systems* **2014**.

[20] S. Weikl, K. Bogenberger, "A practice-ready relocation model for free-floating carsharing systems with electric vehicles – Mesoscopic approach and field trial results", *Transportation Research Part C: Emerging Technologies* **2015**.

[21] Y. Deng, "Design and management of mobility-on-demand (MOD) transportation systems considering demand uncertainty and flexibility-a simulation-based approach", **2015**.

[22] N. Jian, D. Freund, H. M. Wiberg, S. G. Henderson, "Simulation optimization for a large-scale bike-sharing system", *2016 Winter Simulation Conference (WSC)* **2016**.

[23] K. A. Marczuk, H. S. Soh, C. M. Azevedo, D.-H. Lee, E. Frazzoli, "Simulation Framework for Rebalancing of Autonomous Mobility on Demand Systems", *MATEC Web of Conferences* **2016**, DOI 10.1051/matecconf/20168101005.

[24] H. R. Gómez Márquez, R. López Bracho, A. Ramirez-Nafarrate, "A simulation-optimization study of the inventory of a bike-sharing system: The case of Mexico City Ecobici's system", *Case Studies on Transport Policy* **2021**.

[25] Y. Jin, C. Ruiz, H. Liao, "A simulation framework for optimizing bike rebalancing and maintenance in large-scale bike-sharing systems", *Simulation Modelling Practice and Theory* **2022**.

[26] Sid Dhuri, Spatial Data Analysis With Hexagonal Grids, https://medium.com/swlh/spatial-data-analysis-with-hexagonal-grids-961de90a220e, **2020**.