



Distributed Intelligence & Technology  
for Traffic & Mobility Management

# Control strategy for ride-sharing systems by vehicle routing



This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement no. 953783.

Project	DIT4TraM
Grant Agreement No.	953783
Start Date of Project	01-09-2021
Duration of the Project	36 months
Deliverable Number	D3.3
Deliverable Title	Control strategy for ride-sharing systems by vehicle routing
Dissemination Level	Confidential
Deliverable Leader	Université Gustave Eiffel, Lyon, France
Submission Date	28-02-2023
Author	Manon Seppecher
Co-author(s)	Caio Beojone, Minru Wang, Mina Khalesian, Nikolas Geroliminis, Ludovic Leclercq

## Release Approval

Name	Role	Date
Oded Cats	Project reviewer	16-02-2023
Lampros Yfantis	Project reviewer	16-02-2023
Ludovic Leclercq	Task leader	27-02-2023
Nikolas Geroliminis	Work package leader	27-02-2023

## Document History

Version	Description	Name
1	Version for internal review among authors	DIT4TraM_D3.3_v1
2	Version for internal review among project members	DIT4TraM_D3.3_v2
3	Version ready for release	DIT4TraM_D3.3_final

# Contents

<b>1</b>	<b>Introduction</b>	<b>5</b>
<b>2</b>	<b>Related works</b>	<b>7</b>
2.1	Centralized approaches . . . . .	7
2.2	Decentralized approaches . . . . .	8
2.2.1	Auctioning for dispatching and rebalancing . . . . .	8
2.2.2	Driver incentives . . . . .	9
2.2.3	Customer incentives and pricing . . . . .	9
<b>3</b>	<b>Demand estimation for mobility services</b>	<b>10</b>
<b>4</b>	<b>Rebalancing strategies</b>	<b>14</b>
4.1	An auctioning-based approach for ride-hailing fleet rebalancing under uncertainty (UGE) . . . . .	15
4.1.1	Introduction . . . . .	15
4.1.2	Main assumptions . . . . .	16
4.1.3	Vehicle-to-Infrastructure auctioning scheme . . . . .	16
4.1.4	Relocation offers . . . . .	17
4.1.5	Utility functions . . . . .	20
4.1.6	Summary . . . . .	22
4.2	Multi-layer vehicle repositioning (EPFL) . . . . .	23
4.2.1	Main assumptions . . . . .	23
4.2.2	Upper-layer controller . . . . .	24
4.2.3	Middle-layer controller . . . . .	25
4.2.4	Lower-layer controller . . . . .	26
4.3	Rebalancing with in-service vehicles through incentives for ride-splitting (EPFL) . . . . .	28
4.3.1	Main assumptions . . . . .	29
4.3.2	Matching description for pool trips . . . . .	30
4.3.3	Pricing and passenger choice attributes . . . . .	30
<b>5</b>	<b>Methods comparison</b>	<b>32</b>
5.1	Case study description . . . . .	32
5.1.1	Introduction . . . . .	32
5.1.2	Demand balance . . . . .	33
5.1.3	Network specification . . . . .	33
5.1.4	Request generation . . . . .	34
5.1.5	Network partitioning . . . . .	35
5.1.6	Simulation settings . . . . .	35
5.2	Comparative results . . . . .	37
5.2.1	Selected KPIs . . . . .	37
5.2.2	Algorithms comparison . . . . .	38

5.2.3 Discussion . . . . .	45
<b>6 Conclusion</b>	<b>50</b>
<b>Publications for D3.3</b>	<b>52</b>
<b>References</b>	<b>53</b>

# 1 Introduction

Over the last decade, on-demand mobility services have drastically transformed urban transportation ecosystems, driven by the rise of new technologies, especially mobile networks, and terminals. The offer has multiplied, with many Transportation Network Companies (TNCs) appearing on the market, competing with traditional taxi companies and providing travelers with a vast range of services. These services can meet an increasingly dynamic and irregular mobility demand, unsatisfied by public transportation or personal car constraints. On the one hand, ride-sourcing services offer more flexible services than public transit, on-spot and on-demand pickup, and no connections. On the other hand, it can be less costly than private car ownership and provides satisfying solutions to parking issues.

These individual benefits extend to a collective scale, at which the limitation of car ownership can relieve the pressure on land and the need for parking spaces. By encouraging pooled trips through ride-splitting, these services can reduce road traffic and its externalities: participation in urban congestion, air emissions, or noise pollution. Yet, effective management of this type of service requires successfully handling several operational issues: demand prediction, ride-pricing, passenger-vehicle matching, vehicle routing, and fleet rebalancing. The present document discusses this last issue.

Urban mobility dynamics and asymmetry are some of the major issues in managing on-demand mobility services. Without fleet rebalancing strategies, vehicles accumulate in the most attractive areas of the network, to the detriment of the high-demand regions. The satisfaction of new requests may require vehicles to travel long distances to pick them up, implying substantial waiting times for the users. These waiting times affect the attractiveness and performance of the service.

On the contrary, anticipating future demand and relocating the fleet accordingly ensures high service levels to users and minimizes abandonment rates. Thus, two questions arise: how to predict ride-sourcing demand and how to reorganize the fleets accordingly? This paper addresses the latter issue, but we dedicate a preliminary part of our discussion to the former by introducing research work by Université Gustave Eiffel within the DIT4TraM project on the specific topic of demand for mobility services. The method and results briefly presented serve as a methodological foundation for some fleet management strategies described later in the report.

While rebalancing decision-making has long been explored using centralized approaches, the research teams at EPFL, Lausanne, Switzerland, and Univ. Gustave Eiffel, Lyon, France, have developed several distributed rebalancing strategies described and compared in this document. The methods developed cover a large range of methodological approaches. They include:

- An auctioning-based approach to dispatch idle vehicles based on the outcome of a distributed two-sided matching process;

- A hierarchical distributed control strategy to reorganize fleets at the microscopic and macroscopic scales;
- A pricing scheme to rebalancing vehicles through targeted ride-splitting.

A substantial effort was made to build a joint case study to simulate these strategies for the first time in the same environment. The city of Lyon, in France, was selected as a common test case.

This joint analysis evidences a general increase in the number of passengers served, while some of the proposed strategies also reduce passengers waiting time. These strategies also generally contribute to increasing the average service distance of vehicles and decreasing their idle distances. This comparison also allowed identifying the weaknesses of the methods. We discuss them carefully and propose different directions to address them.

This document is organized as follows:

- Section 2 presents the fleet rebalancing related literature;
- Section 3 introduces research conducted on demand estimation for mobility services, developed as a methodological foundation for the development of sound rebalancing strategies;
- Section 4 presents the three rebalancing strategies proposed;
- Section 5 presents the comparison of the rebalancing strategies. It describes the case study and presents and discusses the performances of each strategy.
- Section 6 wraps up this document with a discussion of the results achieved and future research conducted as part of this task.

## 2 Related works

### 2.1 Centralized approaches

Many articles on fleet management in general (routing, dispatching and rebalancing) of ride-sourcing fleets have addressed these issues via centralized methods. The objective of optimal vehicle dispatching is to serve known (present or future, i.e. pre-booked) requests in a way that satisfies an objective, i.e. minimizing the total distances traveled, the costs, the waiting time of passengers, the number of unsatisfied requests or maximizing the number of shared-trips in the case of ride-splitting.

Several rebalancing strategies emerge from the optimization of passenger-driver matching algorithms. These strategies include [1], [2], and [3] who sent empty vehicles to the location of recently unsatisfied customers. However, the issue with these strategies is their reactive nature, in the sense that they mainly accounted for past events (i.e., lost demand) or current conditions. For instance, if an area faces recurrent losses of requests, customers will likely change their travel option to a more reliable transportation mode.

More recent studies try to take centralized actions before losing these passengers. By introducing demand anticipation, the search for the optimal dispatching gradually turns to proactive rebalancing problems. Model Predictive Controllers (MPCs) are then either used to define where each vacant vehicle should relocate [4], or how many vehicles should relocate from a region  $i$  to a region  $j$  [5]. In ride-splitting context, [6] build on the notion of feasible shared trips and leverage future demand prediction in vehicle routing and vehicle assignment. [7] uses coverage control to proactively position idle drivers in areas more likely to originate new requests. [8] uses an MPC to relocate idle taxis in a macroscopic set of regions. However, although highly optimized and proactive, these approaches assume complete compliance to the provided instructions. Hence, they ignore that the objectives of individual humans offering rides, or they assume that the TNC owns a completely autonomous fleet. These strategies also make the systems using them vulnerable to failures or communication interruptions, and rigid in the face of possible competition from other mobility services.

## 2.2 Decentralized approaches

### 2.2.1 Auctioning for dispatching and rebalancing

An early work on the distributed management of ride-splitting fleet with a demand-initiated negotiation scheme is the work of [9]. The authors propose a peer-to-peer ride sharing system in which several types of agents are considered: clients can be mobile or immobile, while vehicles can be either private cars, taxis, or mass transit vehicles. Agents communicate with each other within a limited distance range. The customers send trip requests, to which hosts can respond by making service offers. Customers then select the best offer among the offers proposed.

More recently, a similar approach was developed in [10]. In that framework, if vehicles have free seats, they answer to ride requests by notifying the passenger with the offered route, detour and travel time. Otherwise, they transfer the request to other near-by drivers in a peer-to-peer manner. Again, based on the different offers from the drivers, the rider picks the driver that optimize a multi-objective function.

In [11], the authors propose another agent-based approach for modeling the ride-splitting problem, in which the negotiation process is initiated by the vehicles with vacant seats. Instead of considering potential passengers in their close proximity, they consider and interact with new passengers located along their planned routes. The authors explore different negotiation scheme settings: one-driver-one-rider, one-driver-multi-riders, and the multi-hop ride-splitting problem. Still, the negotiation process is similar to the previous ones, with drivers placing bids on the least expensive ride request, and riders accept the bid with the lowest passenger cost.

Generally speaking, the methods discussed above, although very constructive, cannot be proactive because they involve the passenger in the matching and pick-up repositioning process. As long as the passenger has not appeared, the vehicle remains idle without relocating directions. Therefore, to integrate proactive rebalancing in a negotiation-based fleet management requires either:

1. agents to have intrinsic relocating strategies;
2. or the use of a third-party infrastructure agent.

In [12], authors explored the first option and developed a method for decentralizing taxi decisions in case of a central service shut-down. The second option has, to the best of our knowledge, not been sufficiently investigated yet. While distributed traffic management relies extensively on local controller agents (traffic lights and intersection controllers), few studies have considered infrastructure agents in the ride-hailing and ride-splitting management problems. [13] propose a decentralized ride-splitting matching system based on vehicle-to-infrastructure and infrastructure-to-infrastructure communication scheme, with infrastructure being located at road intersections. This approach still does not address fleet re-balancing

and raises concerns regarding the possibility to scale the method to large road networks over global urban areas, which may contain thousands of intersections.

## 2.2.2 Driver incentives

In the literature, a few efforts acknowledge the impossibility of ride-sourcing service operators deploying or dispatching extra vehicles whenever there is a shortage. In reality, they must be persuasive enough to convince the available pool of drivers to relocate. The works of Lu et al. [14], Sadeghi and Smith [15], and Powell et al. [16] incentivized drivers to decide on the best location for the next assignment through mechanisms that control the supply of drivers and the demand of passengers, such as surge pricing [17]. These strategies also suffer from a reactive nature. Drivers face high uncertainty whether an area previously marked as ‘high demand’ or ‘surge price multiplier’ will be crowded with incoming vacant vehicles, resulting in longer cruising times. Such a problem also emerges because these strategies only provide limited information for drivers (in terms of time and scope), in which case they must estimate themselves what the most profitable option is.

## 2.2.3 Customer incentives and pricing

Ride-sourcing systems benefit further from ongoing developments in the shared economy. Besides driver-based relocation strategies discussed in previous section, where vehicle relocation is carried out by staff and drivers, one can also identify research directions that develop user-based pricing strategies. Earlier research efforts on users’ origin and destination flexibility [18], user-based relocation in car-sharing systems [19], and surge-pricing for ride-sourcing systems [17] demonstrate the potential of tapping into users flexibility in their travel options, as a way to achieve operational goals such as more balanced supply, increased profit, and more attractive service level overall for passengers. Although both aimed at equilibrating demand and supply, surge pricing differs from incentives due to the former’s commonly profit-maximization objectives; customers whose willingness to pay is lower are dissuaded from using the ride-sourcing platform, and adopt alternative travel modes instead. In contrast, the objectives of user-based incentives are commonly service-driven rather than revenue-oriented. From the customers’ perspective, if the lowered price makes the platform appear more attractive without incurring excessive detour, the platform can likely retain more customers compared to the surge pricing case. However, the incentives need to be well-designed to prevent platform revenue loss and account for the uncertainty in user preferences.

### 3 Demand estimation for mobility services

Before presenting developments on fleet rebalancing strategies, this section discusses demand forecasting, a critical dimension for proactive management of ride-sourcing services. It briefly presents recent studies on demand forecasting, then reviews a method designed within the DIT4TraM project to predict demand for mobility services [20]. The methodological features and results are summarized and presented at the end of this section as the approach supports the methodological foundation of specific strategies presented in Section 4.

Mobility services require accurate demand prediction both in space and time to effectively accomplish fleet rebalancing, provide efficient on-demand transportation services, and manage advanced ride-splitting with minimum fleet sizes. Although the optimization of mobility services is a widely studied topic, the critical component of demand prediction has received less attention. New mobility service operations impose new challenges for demand prediction, as high resolution is needed both in space - at large-scale - and in time - at short-term to mid-term (e.g., next fifteen minutes or next hours) - to effectively perform fleet sizing and rebalancing.

Predicting traffic demand throughout a city can help car-sharing companies pre-allocate more cars in high-demand regions or help taxi centers to manage floating taxis by incentivizing vacant vehicles to move from the over-supply regions (the zones with more potential vehicles) to over-demand ones (the zones with more potential passengers) in advance. The same demand-supply imbalance exists in mobility-on-demand services such as e-hailing taxis, which have gained great popularity in recent years. Traffic demand forecasting can help to dispatch cars efficiently and consequently minimize the waiting time for both passengers and drivers [23], [21], [22].

Traffic demand data varies with time and space and has complicated spatial-temporal dependencies. Regarding time dependency, the traffic demand is expected to be high during peak hours (morning and evening peaks) and low at night (sleeping hours). Furthermore, the traffic in each zone depends not only on the historical traffic of that zone but also on the traffic of all the other zones in the whole area of interest, with a stronger impact from nearby zones than distant ones [24], [25].

Traffic demand prediction approaches can be divided into three categories. The first category is the statistical methods [27], [26]. Historical Average (HA), Auto-Regressive Integrated Moving Average (ARIMA), and Vector Auto-Regressive (VAR) are the most well-known statistical methods found in the literature. These algorithms are easy to be deployed but only applicable to relatively small data sets, and the capability of these approaches to deal with complex and dynamic traffic demand data is limited [28]. Traditional machine learning methods [32], [29]–[31] constitute the second category. These methods, such as Support Vector Regression (SVR) and Random Forest Regression (RFR), can process high dimensional traffic

data and capture non-linear relationships.

However, with the advent of deep learning methods [33], [34], [24], [22], which comprise the third category, the full potential of Artificial Intelligence (AI) has been utilized in traffic-related prediction applications [35]. Several deep learning architectures, such as Convolutional Neural Network (CNN), Graph Convolutional Network (GCN) [36], Recurrent Neural Network (RNN) [37], and its variants like Long Short-Term Memory (LSTM) [38] or Gated Recurrent Unit (GRU) [39], have been used for traffic prediction – see [28] for a more extensive survey. In the literature, there are some approaches that have used deep learning for traffic demand predictions [23], [21], [24], [22], [40].

In [40] a feature-level data fusion model is investigated, which integrated the feature attributes into long short-term memory (LSTM), to predict day-to-day travel demand variations based on the origin-destination matrices data of 30 consecutive days on a large-scale transportation network. [22] proposed a sequence learning model based on LSTM for predicting taxi demands and showed that this approach outperformed the feed-forward neural network and naive statistic average predictor. In [24], a deep learning (DL) approach named fusion convolutional long short-term memory network (FCL-Net) was employed by the fusion of convolutional techniques and LSTM network to predict short-term passenger demand for an on-demand transport service. [23] proposed a multi-task deep learning (MTDL) model based on LSTM to forecast short-term taxi demands at a multi-zone level. The proposed model was able to predict the demand of multiple zones simultaneously in a way that the demand prediction of each zone can be conducted by considering the information of zones that can help to improve the prediction. Traffic demand characteristics change at different levels of temporal and spatial aggregations. The aggregation level in the temporal or spatial dimension could affect the prediction accuracy so that aggregating in a longer time window or in more extensive areas results in smoother time series and therefore facilitates pattern identification. However, aggregating in a longer time window or in more extensive areas may cause losing some important information at higher frequencies and remarkably diminish the accuracy. Therefore, finding the right balance is noteworthy [42], [44], [41], [43] and appropriate aggregation in the spatial and temporal dimensions can be used to improve prediction accuracy.

To tackle this problem, we propose an approach that uses hierarchical time series and reconciliation concepts in combination with a deep learning method. A hierarchical time series (HTS) is a collection of time series organized in a hierarchical aggregation structure. Forecast reconciliation is the process of improving prediction accuracy by adjusting the forecasts to make them coherent across the hierarchy. This coherence can be checked at the spatial or temporal levels in time-series data related to traffic demand. For example, at the temporal level, the traffic demands at every hour of a day should add up coherently to give the diurnal traffic demand, or at the spatial level, the forecasts of traffic demand in initial traffic zones (here Traffic Analysis Zones (TAZ)) should add up to provide the forecasts of more extensive areas constructed based on initial traffic zones (here the zones designed to meet the requirements of mobility service operations).

In recent years, hierarchical time series and reconciliation have attracted much attention [45], [46]. However, aggregation in spatial and temporal dimensions has not been considered widely in the transportation field. [42] used a component-wise gradient boosting procedure (CWGB) combined with hierarchical reconciliation to predict traffic flow. In our research, an LSTM approach in combination with hierarchical reconciliation (HR) is proposed for short/mid-term traffic demand forecasting. The proposed approach has three main features that make it different from the previous studies:

- The demand prediction for a specific area can be adjusted based on the demand forecasting of a group of regions at different aggregation levels (spatial levels) and the forecasts in different time resolutions (temporal levels). The main advantage of this feature is that it simultaneously provides the possibility of forecasting demand in different regions with different dimensions and travel times.
- The approach is based on a deep learning technique, here LSTM, combined with an innovative application of HR. Based on this feature, HR improves the initial forecasts resulting from the deep learning method in such a way that the predicted demands for all zones are coherent and consistent with each other.
- The deep learning technique and the HR step are linked together via an error analysis, which can control the solutions to be in a true feasible space and can finally provide us with the expected precision of the forecasted demands. This feature, by providing the prerequisites of hierarchical reconciliation, allows it to be implemented in the most optimal way and the behavior of demand errors can be controlled more accurately.

In this research, we consider two specific spatial partitioning levels: the initial one consists of homogeneously populated areas based on census data, and the second one is the aggregation of these zones into larger areas that still fulfill a maximum service time criterion for mobility services. In short, the later areas should be large enough to allow robust demand predictions but small enough to allow pre-positioned vehicles within the area to serve any internal requests with a low waiting time for passengers. In the proposed method, apart from adopting a deep learning technique for demand prediction, the main idea is that all the predicted demands in all zones and regions of the mentioned partitioning levels should be adjusted and matched with each other in the hierarchical tree structure that is constructed based on these partitioning levels.

In brief, the proposed approach is comprised of three main stages. The aims and ideas of these stages are as follows:

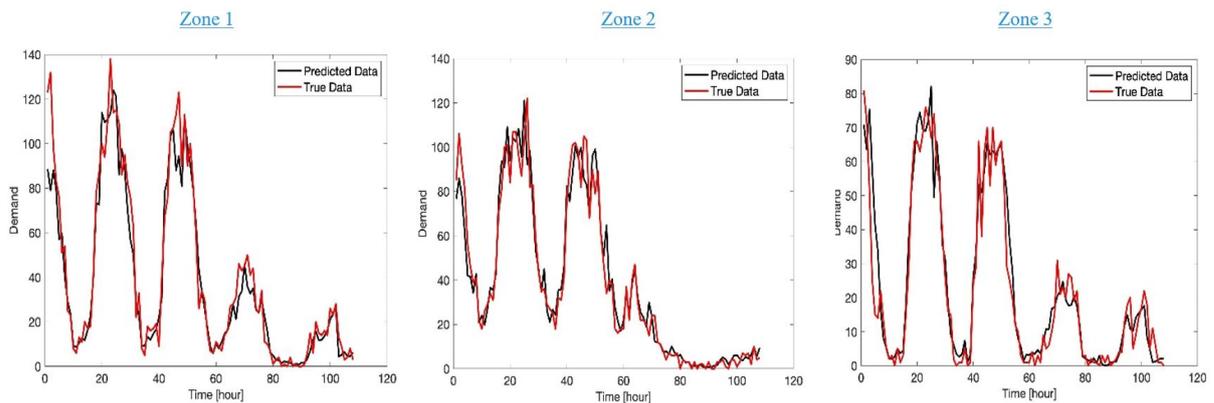
1. Deep learning: This step of the proposed approach provides initial demand predictions for the spatial partitioning levels. In this step, in addition to historical demand data, utilizing other information about the demand as additional input (e.g., external features of time, i.e., time-of-day and day-of-

week) can be beneficial to network training.

2. Error analysis: The main idea behind this step is to validate and control the initial forecasts provided by the deep learning step. The initial forecasts cannot be directly used by the hierarchical reconciliation, and their random errors need to be investigated and controlled. One of the main outputs of the error analysis is the weight matrix, which is essentially required by optimal reconciliation. Meanwhile, the error analysis step leads to the projection matrix, which provides the proposed approach with a filtering. By means of the projection to the feasible space of solutions, the filtering avoids the initial forecasted time series diverging from the true solutions.
3. Optimal reconciliation: the main idea behind this step is to optimally adjust and match all the demands with the hierarchical tree structure of zones. This step leads to coherent forecasts of the demands in all levels of the Hierarchical structure.

We evaluate the proposed approach on a large-scale GPS tracking dataset of Lyon in France. The proposed method reduces the root mean square error (RMSE) by 13.92% and 14.77% for the predefined and aggregated zones, respectively, compared with the LSTM using the historical demand and the external features of time (time-of-day and day-of-week) at fifteen minutes time resolution. Similarly, the corresponding improvement for mean absolute percentage error (MAPE) is 14.87% and 19.23%, respectively.

Furthermore, as an example, in Fig.1 the forecasted values of the demand, obtained by the proposed method, and the true values of the demand have been illustrated for three aggregated zones with the highest values of demand. This figure shows how the results obtained by the proposed approach are in a good agreement with the ground-truth values.



*Figure 1: The forecasted demands and true ones of three zones possessing the highest values of demand amongst the other zones, for the time step equal one hour*

## 4 Rebalancing strategies

We propose three different fleet rebalancing strategies.

The first strategy (M1, Section 4.1) relies on an auctioning-based approach and the rebalancing of vehicles according to the outcome of a distributed two-sided matching process. It is based on the division of the network into service areas supervised by controllers interacting with the mobility service vehicles. It has been designed to allow for the management of different fleets of competing mobility services, and to integrate different types of uncertainties, concerning demand, but also traffic conditions and vehicle travel times. To date, this method has only been applied in a ride-hailing context, but it has been designed to be extendable to ride-splitting operation settings.

The second strategy (M2, Section 4.2) proposes a hierarchical control of the relocation of idle vehicles. At the upper layer, an aggregated model and a model predictive control framework are used to assess the number of vacant vehicles to relocate from one region to another. At a finer scale, a coverage control scheme aims at distributing vehicles according to the demand density within each region. This approach helps bridging the gap between macroscopic and proactive fleet rebalancing and microscopic vehicle management, providing an alternative to distributing the rebalancing task to each area of a city. The method is adapted to ride-hailing and ride-splitting operating modes, and both are discussed in the result section (Section 5).

The final strategy (M3, Section 4.3) uses ride-splitting as a fleet rebalancing strategy. Instead of having solo-riders and the vehicle they were matched with end up in low demand areas, ride-splitting can be used to have vehicles have their final drop-off in a high demand area. This approach may be an alternative to increasing the size of the fleet that would otherwise be necessary to meet demand. As pooled trips can be considered as less attractive than solo trips, the ride-splitting service can apply monetary incentives on the ride prices to increase their attractiveness.

Together, these approaches explore a large range of different strategies to rebalance the fleet of a ride-sourcing service.

## 4.1 An auctioning-based approach for ride-hailing fleet rebalancing under uncertainty (UGE)

### 4.1.1 Introduction

The first fleet management strategy for on-demand mobility services focuses on distributing empty vehicles on the urban network. We design a decentralized decision architecture consisting of a mesh of controllers that divide the urban network into an equal number of service areas. These agents are considered to be at the service of a public authority (e.g., a transport agency or a local authority) and aim to satisfy their local demand as fast as possible. These controllers can be coupled with physical infrastructures, such as parking lots or depots for vacant vehicles. Note that the work presented in this deliverable assumes that the vehicles belong to a single mobility service. We will explore later a case study including different competing mobility services in connection with WP5.

To guarantee the fast pick-up of the local demand, the controllers are first in charge of predicting the number of travelers requesting a ride within the service area. This prediction can be based on local demand history and known pre-booked requests if any. This prediction allows controllers to estimate the number of vehicles needed to meet local demand and to implement a negotiation process with cars to attract the required number.

These negotiations between vehicles and controllers are done simultaneously and in a decentralized way through a two-sided matching market. Vehicles apply to their favorite relocation offer, *i.e.*, the offer that will maximize their expected revenue, while controllers aim at ensuring the fastest service for the local passengers. This reconciliation may require several iterations of the process, with vehicles that have been rejected by their preferred region applying for the next one. At the end of this process, vacant vehicles are assigned to a service area to which they will relocate. Figure 2 illustrates this communication protocol.

Note that although the management of rebalancing is outsourced from the mobility service to a public authority, we consider that the mobility service keeps the management of the matching and routing of its passengers (not represented in Figure 2). It is assumed at this point that ride-hailing services and vehicles comply with the outcome of this process. Future work could examine how non-compliant services and vehicles can alter system performance.

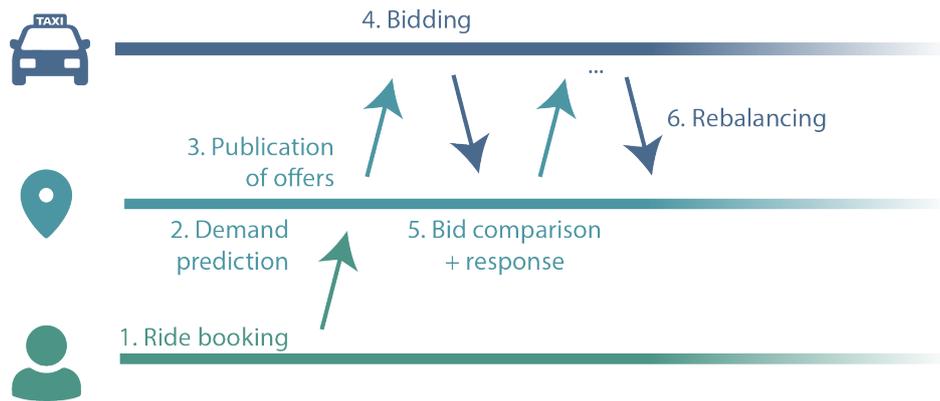


Figure 2: Communication protocol between travelers, controllers, and vehicles

### 4.1.2 Main assumptions

The main assumptions supporting our method are:

1. Local controllers can predict the local demand and its uncertainty based on historical data (supported by results of Section 3), as well as the average zonal revenue;
2. They publish relocation offers with a probability higher than a pre-defined threshold;
3. Vehicles utility is based on expected revenue and rebalancing costs;
4. Controllers utility is based on the vehicles arrival time;
5. The auctioning process is truthful;
6. The mobility services and their fleet comply with the rebalancing directions.

### 4.1.3 Vehicle-to-Infrastructure auctioning scheme

The rebalancing strategy developed at the upper-level relies on an auctioning framework, involving local controllers negotiating with vehicles for the attribution of relocation options. On the one hand, local controllers aim at maximising the satisfaction of their future local demand by minimising their waiting time before pick-up. On the other hand, vehicles aim at relocating towards area that will maximise their net profit, *i.e.* minimise their relocation costs and maximise their service incomes.

The auction framework is based on the **distributed** Gale-Shapley algorithm.

We summarise below the outline of our algorithm:

- Local controllers predict the future local demand (*potential future passengers*), and deduce from it their future vehicle needs. They express those needs by publishing *relocating options* into a two-sided *relocating market*. Each relocat-

ing option is associated with the likelihood that the demand will occur.

- With a fixed frequency, vacant vehicles can enter the market and evaluate for each relocating option their own utility. Vehicles estimate this utility by considering their relocation costs towards the service area, the likelihood of the demand associated with the relocating options, and the incomes expected for serving this demand. On this basis, the vehicles make an ordered list of the relocation options they consider useful, and bid on the first one. When bidding, vehicles declare their expected arrival time at destination.
- Local controllers agents gather the bids of vehicles, and evaluate for each the waiting time imposed on the future passenger. For each potential passenger, regions accept the best vehicle's bid and reject the others.

The auctioning process continues in an iterative way until rejected vehicles run out of bids. This process is further described in Algorithm 1.

#### 4.1.4 Relocation offers

The purpose of the auctioning game between service areas and vehicles joining is to allocate relocation offers, published by service area agents in charge of their prediction. Let  $i$  be a service area,  $T$  a time period, and  $k$  an integer index. We note  $r_{i,T}^k$  the  $k^{\text{th}}$  relocation offer of region  $i$  for period  $T$ , defined as the tuple  $r_{i,T}^k = (p^k, \hat{g}_{i,T}^k)$  with:

- $p^k$  the occurrence probability of the  $k^{\text{th}}$  most probable ride request in  $i$  during  $T$ ;
- $\hat{g}_{i,T}^k$  the expected income from serving the corresponding potential passenger. In this paper, we will assume  $\hat{g}_{i,T}^k$  is independent of  $k$  and  $T$ , therefore  $\hat{g}_{i,T}^k = g_i$ .

Hereafter we describe those relocation offers and their components. For readability, we will hereafter also use  $k$  to designate the relocation offer  $r_{i,T}^k$ .

##### 4.1.4.1 Occurrence probability $p^k$

Let  $i$  denote a service area and  $T$  a prediction period.

The local total demand (or total number of ride requests within  $i$ ) for period  $T$  is modelled as a random variable  $X_{i,T}$ . We assume that this total local demand  $X_{i,T}$  follows a normal distribution:

$$X_{i,T} \sim N(\mu_{i,T}, \sigma_{i,T}^2) \quad (1)$$

with  $\mu_{i,T}$  and  $\sigma_{i,T}$  be predicted from historical analysis of the data.

Relocation offer generation emerges from the breakdown of the demand distribution into demand units, each representing a possible additional customer. These

---

**Algorithm 1: Rebalancing algorithm**

---

```
1 foreach  $i \in \text{service areas}$  do
2   Gather booked rides;
3   Predict and split-up future demand;
4   Publish related rebalancing offers into the matching market;
5 foreach  $v \in \text{idling vehicles}$  do
6    $l \leftarrow []$  foreach  $i \in \text{service area}$  do
7      $c_i \leftarrow \text{compute\_rebalancing\_cost}(v, i)$ ;
8     foreach  $T \in \text{in horizon periods}$  do
9       foreach relocation offer  $r_{i,T}^k$  do
10         $U_v(k) \leftarrow \text{get\_utility}(r_{i,T}^k, c_i)$  (See Section 4.1.5.1)
11         $l \leftarrow [l, U_v(k)]$ 
12    $l \leftarrow \text{descending\_sort}(l)$ ;
13 while  $\exists$  unmatched driver with non empty application list do
14   foreach unmatched driver with non empty application list do
15     Driver applies to most preferred rebalancing option;
16   foreach rebalancing option do
17     Compute bids utilities (See Section 4.1.5.2)
18     Select best bid;
19     if option is not matched yet then
20       Match with vehicle;
21       Reject all others;
22     else
23       if utility of best bid outperform current utility then
24         Unmatch current driver ;
25         Match with best bidding vehicle ;
26         Reject all others;
27       else
28         Reject all applying vehicles;
29   foreach unmatched driver do
30     Remove previous application from application list;
```

---

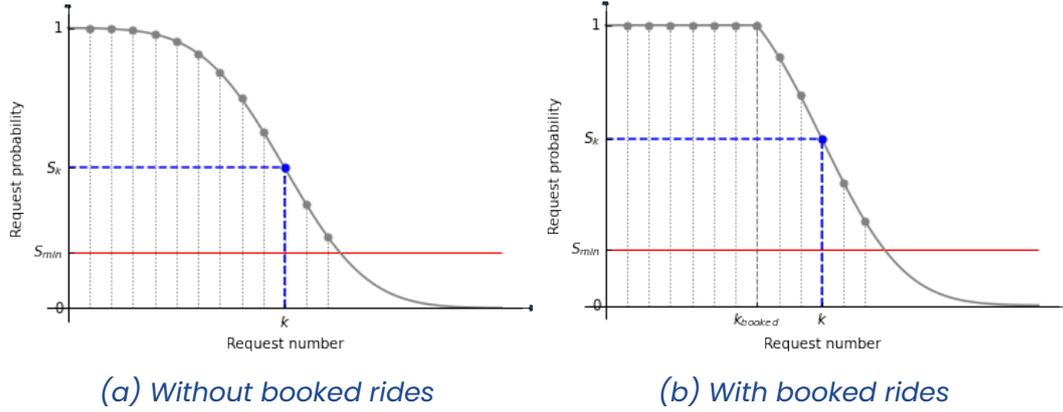


Figure 3: Demand prediction split-up

demand units are characterised by an occurrence probability  $p^k$ , so that:

$$\forall k \in N, \quad p^k = S_{X_{i,T}}(k) = P(X_{i,T} \geq k) \quad (2)$$

where  $S_{X_{i,T}}$  is the survival function of  $X_{i,T}$ . Hence,  $p^k$  describes the probability that at least  $k$  ride requests will occur during  $T$  in the considered service area.

In theory, local controllers can generate as many offers as they wish. But the probability of each additional request is lower than the previous one, and therefore less interesting for bidding vehicles. In practice, we define a probability threshold  $p_0 \in [0, 1]$  and assume each local controller publishes on the relocating market only the  $k_{max}$  relocating offers with occurrence probability above  $p_0$ . Figure 3a illustrates the partitioning of the predicted future demand into uncertain requests.

#### Case of pre-booked rides

Additionally, we consider passengers can book rides in advance. Upon the reception of  $k_{booked}$  pre-booked ride requests for a period  $T$ , the local controller can safely publish on the matching market  $k_{booked}$  certain ride requests. It still needs to assess the likelihood of additional requests appearing in real time knowing that

$$X_{i,T} \geq k_{booked}:$$

$$\forall k \in N, \quad p^k = S_{X_{i,T}|k_{booked}}(k) = P(X_{i,T} \geq k | X_{i,T} \geq k_{booked}) \quad (3)$$

$$= \frac{P(X_{i,T} \geq k \cap X_{i,T} \geq k_{booked})}{P(X_{i,T} \geq k_{booked})} \quad (4)$$

If  $k \leq k_{booked}$ , we retrieve:

$$S_{X_{i,T}|k_{booked}}(k) = \frac{P(X_{i,T} \geq k_{booked})}{P(X_{i,T} \geq k_{booked})} = 1 \quad (5)$$

On the contrary, if  $k > k_{booked}$ :

$$S_{X_{i,T}|k_{booked}}(k) = \frac{P(X_{i,T} \geq k)}{P(X_{i,T} \geq k_{booked})} = \frac{S_{X_{i,T}}(k)}{S_{X_{i,T}}(k_{booked})} \quad (6)$$

Then,  $S_{X_{i,T}|k_{booked}}(k)$  can be easily derived from Equation 1. Figure 3b illustrates the demand split-up with booked requests.

#### 4.1.4.2 Income estimation $\hat{g}_i$

Besides request probability, controller agents share to drivers the average income they can expect from serving a passenger from the corresponding service area. This way, vehicles evaluate and compare their utility in relocating towards one region or the other. In a real-world context, this information can easily be derived from historical data. Assuming the ride pricing scheme  $c_{ride}$  is distance-based, as often implemented in the literature, here we build up a simple proxy for historical data to estimate this average local income information. It is based on shortest-path information between regions  $d_{SP}(i, j)$ , ride-pricing function  $c_{ride}$  and standard origin-destination matrix  $M = (m_{i,j})$ :

$$\hat{g}_i = \frac{\sum_j m_{i,j} \cdot c_{ride}(i, j)}{\sum_j m_{i,j}} \quad (7)$$

### 4.1.5 Utility functions

The auctioning process involves that, in turn, vehicles evaluate their utility in applying to one or the other relocation offer, and regional controllers evaluate the utility of these applications for the expected passengers. Below are described the details of those computation process.

#### 4.1.5.1 Utility of vehicles

Vehicles compute their utility  $U_v(i, k)$  in relocating towards region  $i$  to serve potential passenger  $k$  as the net profit of the rebalancing, *i.e.* as the difference between their expected earning  $g(i, k)$  and the relocating costs  $c_{rebalancing}(v, i)$ :

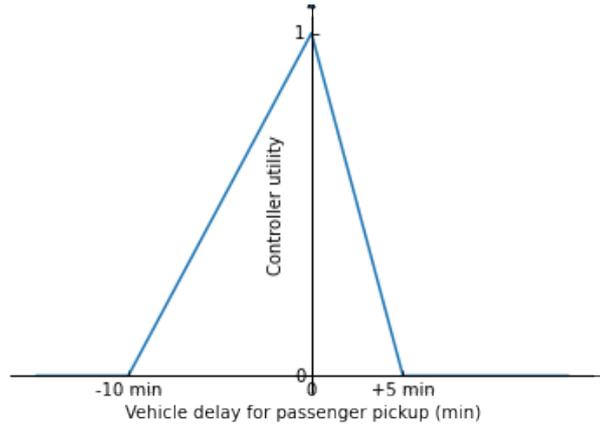


Figure 4: Utility function of local controllers

$$U_v(i, k) = g(i, k) - c_{rebalancing}(v, i) \quad (8)$$

The rebalancing costs are computed as the product of the vehicle' mileage cost  $c_v^{km}$  with the relocating distance from current position  $p$  to region  $i$  following shortest-path  $d_{SP}(p, i)$ :

$$c_{rebalancing}(o, i) = c_v^{km} \cdot d_{SP}(p, i) \quad (9)$$

We take into account the risk that the ride request does not occur. Therefore, the expected gain  $\bar{g}_i$  (cf. Equation 7) is weighted by the probability of occurrence of the query (cf. Equation 2). We define  $g_1(i, k)$  as:

$$g_1(i, k) = \bar{g}_i \cdot S_{X_{i,T}}(k) \quad (10)$$

#### 4.1.5.2 Utility of local controllers

We plan to explore different controllers utility function to identify which vehicles satisfy the best their objective to limit the waiting time of their own customers. These utility function can take into account uncertainty related to the passengers and vehicles arrival time. Yet, to compute the results presented in this document we use a simple utility computation process. Within the rebalancing period, the local controller randomly draws a time at which the predicted passenger is assumed to arrive. Then, each vehicle application is evaluated depending on the predicted delay of the bidding vehicle, as illustrated in Figure 4. This utility function can be evaluated in light of the number of matched passengers. Sensitivity analyses will allow to calibrate it in a more refined way, in order to maximize the number of passengers served.

## 4.1.6 Summary

To summarise, we design an auctioning-based rebalancing framework which relies on the external control of vehicles' macroscopic rebalancing by cross-platform service area agents. At each rebalancing time step, local controllers gather pre-booked rides, predict the overall future demand and publish relocation offers on a two-sided matching market. Idle vehicles bid on their most profitable option, and controller positively or negatively answer depending on the capacity of vehicle to arrive on time to satisfy the anticipated customer. Rejected vehicles apply to their second most profitable option, while regions may reject already accepted offers if they receive a more useful one. The outcome of this matching process determines how vehicles should rebalance over the planning horizon. In-between two relocation periods, idle vehicles are open to matching with new ride requests occurring in their nearby environment.

Future developments of this methodology will include the exploration of more complex utility functions and advanced sensitivity analysis to uncertainty levels and number of relocation published. In the long run, this approach will be used to devise local incentive strategies to encourage vehicles to relocate in service areas suffering from lower accessibility or uncertain demand. We will also explore the performances of this method to manage the rebalancing of multiple competing mobility services in connection with works conducted in WP5. In particular, it may be possible to assess how these third-party controllers can mitigate competition to promote equity among operators of different fleet sizes. We will also look at enriching the method to foster cooperation between local controllers rather than competition.

## 4.2 Multi-layer vehicle repositioning (EPFL)

We propose a hierarchical control strategy for the relocation of idle ride-sourcing vehicles, and in specific for addressing the gap between proactive repositioning strategies and micro-management of vehicles in such activities. The upper layer utilizes an aggregated model, which is an approximation of trip-based MFD modeling approach (building on [47]). A model predictive control (MPC) framework is employed to determine the number of idle vehicles to be relocated for each pair of regions. Unlike perimeter control MPC methods, fleet management MPC requires the integration of more sophisticated MFD-based models describing mixed dynamics of private vehicles and taxis. In the lower layer, given the demand density over the current region, a coverage control scheme operates to distribute the vehicles within the region to achieve a demand-aligned configuration, which provides each vehicle with relatively detailed (i.e., intersection/node-level) position guidance. To bridge both layers, a middle layer mechanism is developed for converting the upper layer decisions into dispatching commands for individual vehicles by solving a Mixed-Integer-Linear-Problem, which minimizes the distance required to achieve the optimal coverage and repositioning decisions.

Implementing a controller for a large-scale system, one may face problems such as high computational effort due to complex models and high dimensions required for accurate network modeling, especially if the model and controller are developed to compute control actions for every individual vehicle over the whole network. One way to solve this problem is to build a hierarchical control structure. Such structures decompose the control problem into a hierarchy of decision-making levels, and operate via coordinating between the actions of an upper layer controller (operating at the aggregated traffic level) and a lower layer controller (managing individual vehicles). The control structure is shown in Figure 5.

### 4.2.1 Main assumptions

The main assumptions supporting this method are:

1. The network can be separated into regions of homogeneous congestion (MFD-based model requires this to run predictions in the upper-layer);
2. The network controller can predict regional demand (as a Poisson process) for the near-future;
3. The local controller can use historical data to determine the intra-regional demand distribution (as a probability function on the origins of new service requests);
4. Vehicles are fully compliant to the provided repositioning and dispatch instructions;

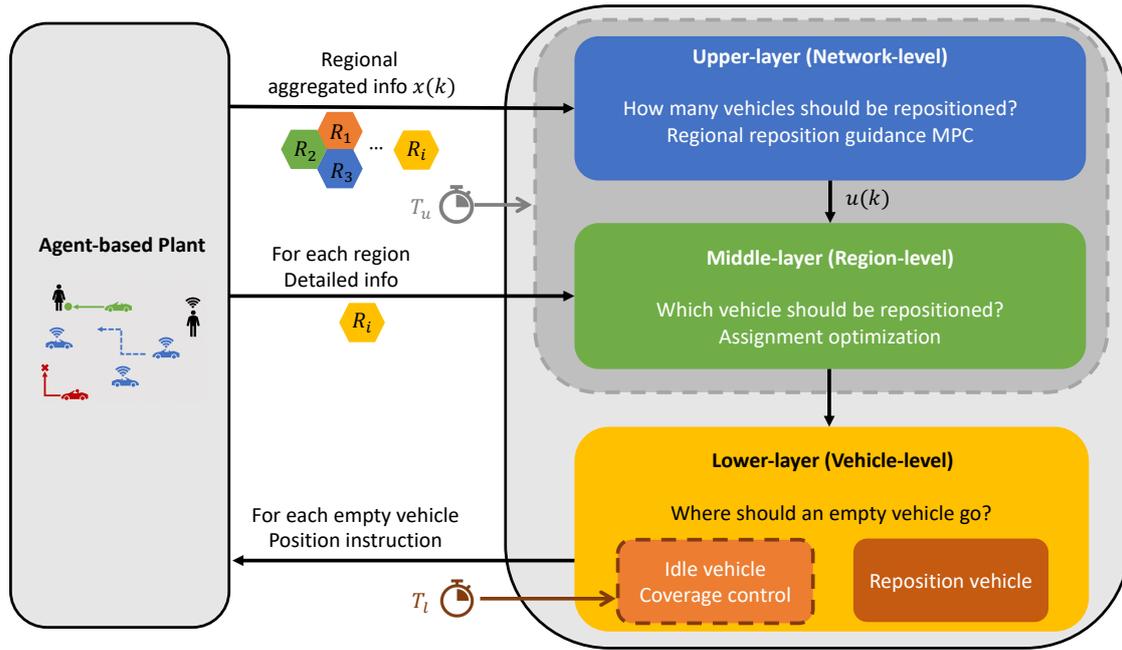


Figure 5: A hierarchical control framework for vehicle rebalancing.

## 4.2.2 Upper-layer controller

The upper layer controller collects aggregated information, such as how many empty vehicles are in each region, from all urban regions at a relatively large update period  $T_u$ . The control action generated from the upper layer determines how many vehicles should stay in current regions and how many vehicles should relocate to other regions, in order to improve availability and thus minimize the total waiting time of passengers. Furthermore, the middle layer transfers the obtained upper layer guidance to the lower layer and specifies which vehicle should stay or move, considering the travel costs caused by repositioning. It is operated within each region and requires relatively more detailed information, such as the coordinates of each vehicle and whether it is occupied or not. Note that the middle layer can only be activated when the upper layer is active. The lower layer is operated in a distributed manner so that each vehicle can obtain its own control action, which facilitates its implementation at a fast update period  $T_l$ . The empty vehicles that are instructed to stay in the current region (i.e., idle vehicles, see the left part of lower layer in Figure 5) communicate and cooperate with each other to achieve better vehicle position configuration, while the rest of the vehicles (i.e., repositioning vehicles, see the right part of the lower layer in Figure 5) are guided to other desired regions as per the relocation commands.

In the upper-layer, dynamic equations can be discretized in time with a sampling time  $T_u$  (unit: h), for enabling formulation of an associated finite-dimensional nonlinear optimization problem. Rewriting them in a compact form, we can formulate the problem of finding the optimal relocation control input  $r_{od}(k)$  values that minimize the total number of canceled trip requests, as the following discrete-

time economic nonlinear MPC problem:

$$u_\kappa \quad \sum_{\kappa=1}^N \sum_{o \in \mathcal{R}} \sum_{d \in \mathcal{R}} \lambda_{od}^B(k + \kappa) \exp(\gamma_0 (n_{od,\kappa}^V)^{\gamma_1} (v_{o,\kappa})^{\gamma_2} w^{\gamma_3}) T_u \quad (11a)$$

$$x_0 = x(k) \quad (11b)$$

$$\text{for } \kappa = 0, \dots, N - 1 : \quad (11c)$$

$$x_{\kappa+1} = F(x_\kappa, q(k + \kappa), u_\kappa) \quad (11d)$$

$$0 \leq Tr_{od,\kappa} \leq n_{od,\kappa}^V \text{ for } o, d \in \mathcal{R} \quad (11e)$$

where  $\kappa$  is the MPC time interval index (i.e., discrete-time clock internal to the MPC),  $k$  is the current discrete time step,  $N$  is the prediction horizon, while  $x_\kappa$  and  $u_\kappa$  are the state and control input vectors internal to the MPC (i.e., predicted states and controls), respectively. To compute the loss probability,  $\lambda_{od}^B(k)$  is an exogenous demand input for busy ride-hailing vehicles; the remaining  $\gamma_i, i \in \{0, 1, 2, 3\}$  are parameters expressing the sensitivity of the matching algorithm to the number of vacant vehicles ( $n_{od,\kappa}^V$ ), traveling speeds ( $v_{o,\kappa}$ ) and passenger waiting time tolerance ( $w$ ). A detailed estimation of parameters in Equation 11 will be provided in the full paper.

### 4.2.3 Middle-layer controller

Once the upper layer provides the number of vehicles transferring between regions, the mid-layer selects which vehicles should move to other desired regions. Vacant vehicles staying in the current region are operated to maintain a good spatial configuration to uphold service quality. Therefore, for each region  $R$ , the coverage control method is employed to compute the optimal configuration of intra-regional positions of individual vehicles for maximizing the weighted covered area. Then, given the computed intra-regional vehicle positions, the assignment of vehicles can be determined by solving the optimization problem in Equation (12), complying with the upper layer decision. Specifically, for objective  $D_R$ , the first term considers the reposition distance between regions, while the second term takes the intra-regional traveling distance into account:

$$\psi_{ir}, \omega_{il} \quad D_R = \sum_{i \in I} \sum_{r \in \mathcal{R}} \psi_{ir} d_{ir}^{\text{out}} + \sum_{i \in I} \sum_{l=1}^{u_{RR}} \omega_{il} d_{il}^{\text{in}} \quad (12a)$$

$$\sum_{r \in \mathcal{R}} \psi_{ir} + \sum_{l=1}^{u_{RR}} \omega_{il} = 1 \quad \forall i \in I \quad (12b)$$

$$\sum_{i \in I} \psi_{ir} = u_{Rr} \quad \forall r \in \mathcal{R} \setminus \{R\} \quad (12c)$$

$$\sum_{i \in I} \omega_{il} = 1 \quad \forall l \in \{1, \dots, u_{RR}\} \quad (12d)$$

$$\sum_{i \in I} \sum_{l=1}^{u_{RR}} \omega_{il} = u_{RR} \quad (12e)$$

$$\psi_{ir} \in \mathbb{B} \quad \forall i \in I \text{ and } \forall r \in \mathcal{R} \quad (12f)$$

$$\omega_{il} \in \mathbb{B} \quad \forall i \in I \text{ and } \forall l \in \{1, \dots, u_{RR}\} \quad (12g)$$

where  $\psi_{ir}$  and  $\omega_{il}$  are binary decision variables expressing whether a vehicle  $i$  is assigned to region  $r$  or to position  $l$ ,  $d_{ir}^{\text{out}}$  and  $d_{il}^{\text{in}}$  are the traveling distances for a vehicle  $i$  to reach region  $r$  or the position  $l$  in the current region, respectively,  $I$  is the set of vacant vehicles,  $u_{Rr}$  is the control action obtained in the upper layer, indicating how many vehicles should move from region  $R$  to region  $r$ , while  $u_{RR}$  is the number of idle vehicles that should stay in the current region.

Note that the constraint in Equation (12b) limits a vehicle either to stay in the current region or move to another region. Equations (12c) and (12e) ensure compliance with the upper layer decision. This problem was inspired by classic assignment problems, such as [48].

#### 4.2.4 Lower-layer controller

In the lower layer, the coverage control algorithm is operated for the vehicles that are instructed to stay in the current region (i.e, idle vehicles). The coverage controller steers these vehicles towards an optimal spatial configuration (indirectly, towards maximizing availability for service) by operating at a fast time scale and with detailed position guidance [49]. Such coordination provides benefits to the system by dynamically allocating the vehicles according to the different demand densities of various city districts.

The city map can be presented as an undirected graph  $G = (Q, E)$ , where  $Q$  is the set of nodes representing the intersections and  $E$  is the set of road links. If origin-destination pairs for trips are recorded in historical taxi data, we can compute the probability that a request starts at a node as  $\phi(q)$ . With a slight abuse of notation,  $q$  in this section denotes a node on the graph (with  $\sum_{q \in Q} \phi(q) = 1$ ).

The coverage objective function in Equation 13 is formulated using algorithms that

group the nearest intersections to each vehicle. This information allows us to define the so-called Voronoi tessellation. According to [50], the optimal position configuration of all vehicles is attained when each vehicle is at the centroid of its respective Voronoi cell. The centroid of a graph Voronoi cell can be computed as an integer optimization problem as Equation 14. As it only requires local information for each vehicle to calculate the Voronoi tessellation, this control algorithm is able to provide each vehicle with an intersection/node-level rebalancing command in a distributed manner (i.e., without requiring a central planner to coordinate the movements of all vehicles).

$$H(P, V) = \sum_{i=1}^{n_{idle}} \sum_{q \in V_i} d(p_i, q)^2 \phi(q) \quad (13)$$

$$C(V_i) = \arg \min_q \sum_{q \in V_i} d(p_i, q)^2 \phi(q) \quad (14)$$

### 4.3 Rebalancing with in-service vehicles through incentives for ride-splitting (EPFL)

Given a ride-sourcing platform that dispatches a fleet of vehicles to serve single-occupancy solo trips and double-occupancy pool trips, the third proposed rebalancing strategy utilizes pool trips as a rebalancing tool. The operator can employ this strategy by extracting information on demand density and demand loss distributed in different areas of the network, and recommend vehicles to serve a specific set of pool trips, so that these drivers can drop off their final passenger in a high demand area, whereas without the rebalancing strategy vehicles are likely to end up in low demand areas.

The benefit of rebalancing in-service vehicles is as follows: given some spatially heterogeneous ride-sourcing demand which cannot be served in a satisfactory manner without increasing the fleet size, for example, as illustrated in [51], one can observe a shortage of vacant vehicles in high demand regions in the network. This means that empty vehicles often immediately become assigned to pick up some new passengers in their vicinity, and may not remain empty for a long time for relevant rebalancing strategies to be applied. Therefore, idle vehicle relocation strategies such as [1] alone may no longer be sufficient in improving the service level. Additionally, we expect rebalancing of in-service vehicles to reduce the total distance that vehicles must travel for the case of empty vehicle rebalancing.

As an initial step, all passengers are assumed to have the same preferences and willingness to pool in response to some trip attributes, which will be elaborated below in Equation (15). Generally speaking, longer trip duration, higher trip cost, and the presence of another rider, will be considered as undesirable when multiple options are available. Therefore, when the ride-sourcing platform proposes a pool trip as a means of rebalancing and passengers have the choice to accept or reject the pool trip, an appropriate monetary incentive can be offered in order to increase the attractiveness of the pool option, so that the operator can better meet its rebalancing objective. The assumption on passengers' choice preference being homogeneous can be relaxed at a later stage.

A user-based incentive naturally increases the attractiveness of a travel option for the passenger, when other trip attributes remain unaltered; however, without any policy for subsidies from the platform, the pricing incentive is directly related to a revenue reduction for the ride-sourcing platform. The methodology developed later in this chapter will consider a baseline pricing strategy as established by Uber, which is regarded as a cost-neutral policy. The second proposed rebalancing strategy without pricing incentive is also considered to be cost-neutral for the operator, since no additional discounts are provided for the rebalancing trips. The third proposed rebalancing strategy offers an additional \$1-off to pool trips that can serve as rebalancing trips. If the number of rebalancing trips increases, this strategy may eventually lead to a platform revenue reduction; therefore, a next research direction should be to consider any trade-off between matching performance and the

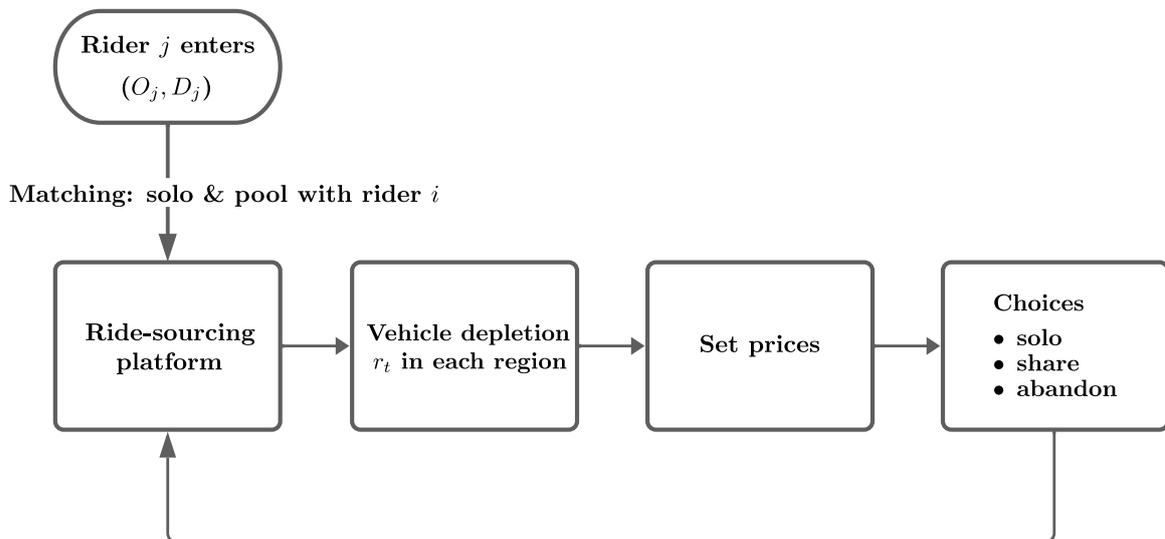


Figure 6: Simulation framework for in-vehicle rebalancing with pool incentives

costs of incentives.

The rebalancing framework is implemented in a discrete event simulator illustrated in Figure 6, which describes the workflow every time a new request from rider  $j$  enters the platform. First, a maximum of two matching vehicles are proposed, with one being a solo option and second a pool option with rider  $i$ . For the pool option, only the vehicle as the best pool candidate is selected. This refers to the vehicle that, when a pool trip is served, would lead to the lowest ratio between the total distance of the pool alternative and the total solo distances if served solo. Then, for the origins and destinations of riders  $i$  and  $j$ , the platform evaluates the current vehicle demand in each region and classifies the regions as high or low demand areas. Next, by comparing the demand characteristics at the origins and destinations of riders  $i$  and  $j$ , we determine whether the pool option can serve as a rebalancing trip, and whether solo option will impede rebalancing actions. Finally, the rider(s) choices between solo and pool options are characterized by a random draw from a probability distribution. The remainder of this section will provide details on the above four components.

### 4.3.1 Main assumptions

This methodology relies on the following main assumptions, with most similar to those stated for Method 2 presented in Section 4.2.1:

1. One single MFD can capture the congestion dynamics of the network for travel time estimation.
2. Private vehicles circulate in the network to travel from their origin to destina-

tions; they are assumed to be parked before and after completing their trips, and therefore do not contribute to congestion.

3. Passengers have identical choice preferences given the attributes of several travel options.
4. Vehicles are compliant when given pick up and drop off instructions. In particular, they are indifferent to any revenue estimation.

### 4.3.2 Matching description for pool trips

When a new rider, referred to rider  $j$ , places a request, a solo trip is proposed to match the rider with the nearest empty vehicle, provided that matching constraints are satisfied. Additionally, a pooled trip is proposed. As summarized in Table 1, parameter  $\Psi$  is set such that the total distance of the pool trip must not exceed the sum of the lengths of two respective solo trips. The platform computes the ratio between the first and second quantity. In the base pooling scenario, the pooling option with the lowest ratio is proposed.

In the rebalancing scenario, given the respective origins and destinations of potential pooling, the platform identifies the respective zones as exhibiting high or low demand to assess the necessity of rebalancing actions. A rebalancing pool trip is loosely defined as one where the vehicle has its final destination in a zone with high demand loss, whereas if two separate solo trips were to be served, one vehicle would drop off its passenger in a low demand zone, hence contributing to the accumulation of empty vehicles in the area, unable to fulfil requests in high demand areas that are far away. For this purpose, the proposed method works with a network where nodes are partitioned with a k-means clustering algorithm into a predefined number of regions that each contains demand aggregated at the region centroid, as described earlier in Section 5.1.5. This level of aggregation can help observe spatial and temporal demand patterns.

### 4.3.3 Pricing and passenger choice attributes

The next step consists of setting the respective prices for the solo and pool trip. First, a baseline price for each service is proposed, which is composed of a base price plus a per-kilometer price, as presented in Table 1. Then, when a pool trip is considered to be desirable as a rebalancing trip, the platform suggests this trip to the rider. Two sets of experiments are run: one where the platform provides no pricing incentive, and another where a flat-rate discount of \$1.00 is proposed to each rider in a pool trip. As discussed in the literature review in Section 2, in user-based relocation strategies, the platform can increase the attractiveness of this option for the passengers, as a way to encourage them to choose one option over another, and help achieve the platform's rebalancing objectives.

When traveller  $i$  is presented with one solo and one pool trip option, their choice

probability is related to their perceived costs associated with each option. The proposed model draws a random number  $x_h^i$  with uniform probability between 0 and 1, and compare this value with the perceived attractiveness of the pooling option  $\bar{h}^i$ , which represents the pooling attractiveness  $h_i$  normalized with an estimated lower and upper bound on its parameters. Define the attractiveness of pooling  $h^i$  for a rider  $i$  as

$$h^i = \alpha \cdot \nu^i \cdot (t_{travel}^{i, solo} + t_{wait}^{i, solo} - t_{travel}^{i, pool} + t_{wait}^{i, pool}) + c_{pool} + (p^{i, solo} - p^{i, pool}) \quad (15)$$

where, from left to right of Equation (15),  $\alpha$  is a parameter tuned based on the lower and upper bounds of  $h^i$ ,  $\nu^i$  is the passenger's value of time in \$/hour; the next four terms are the travel and waiting times of the solo and pool options, respectively;  $c_{pool}$  is another tuned parameter that reflects the relative comfort of a pool trip compared to a solo one, and is identical across the population; finally, the final two terms together represent the price difference between the solo and pool option. As for the normalization of  $h^i$  for obtaining  $\bar{h}^i$ , this is performed by considering the lower and upper bounds for travel time, waiting time, and price, i.e. information related to the minimum/maximum trip lengths, waiting times, and prices. This ensures that the normalized values are mainly between 0 and 1. If  $x_h^i \leq \bar{h}^i$  holds, the rider chooses pool, and otherwise solo; in other words, a higher  $\bar{h}^i$  is related to a higher probability of acceptance for pooling. Both riders must accept pooling for a match to be successful, i.e. mutually accepted.

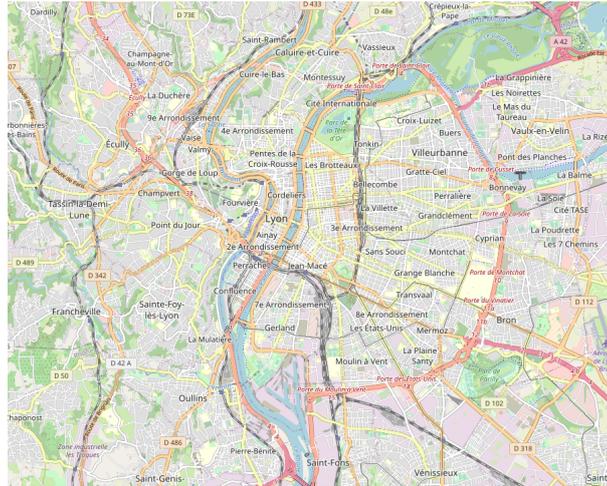


Figure 7: OSM map of Lyon

## 5 Methods comparison

### 5.1 Case study description

#### 5.1.1 Introduction

We choose the city of Lyon, France as a case study. Lyon is the third most populated city in France and the center of the second most populated urban area in France.

The network we model covers 121 km<sup>2</sup> and includes both the city of Lyon and the city of Villeurbanne, both located within a circular ring road. The city is organized at the confluence of two rivers, the Rhône and the Saône Rivers. The city land between the two rivers, hosts the commercial city center. The western part of the city includes the historic district, while the eastern half of the city hosts the universities, business areas and major transportation hubs such as Lyon Part-Dieu train station and Saint-Exupery Airport. Figure 7 presents a map of the city together with its main transportation infrastructures.

The supply calibration and the demand scenarios used in here have been calculated within the ERC Magnum Project [52].

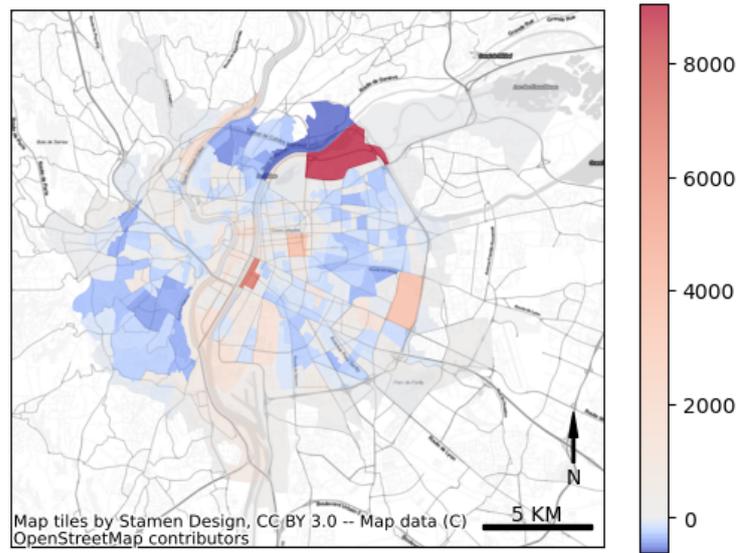


Figure 8: Travelers inflow and outflow during the morning peak

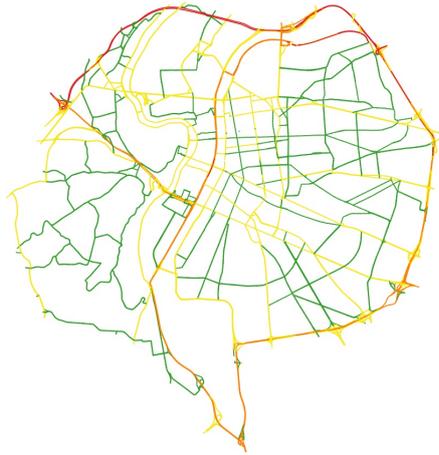
### 5.1.2 Demand balance

In Figure 8, we present the outgoing demand deficit over the study area during the morning peak (6am to 10am). This representation allows to identify the geographic patterns of the demand during that period, and anticipate the needs for fleet rebalancing. We observe that some specific areas are characterized by large imbalances of incoming and outgoing flows, with high inflows and low outflows. It is especially the case for university areas: La Doua campus in the north of the city and Université Lyon 3 campus along the Rhône river. The area of the train station Lyon Part-Dieu, in the city center, and the Vinatier hospital area, in the eastern part of Lyon also seem equally attractive. Although these spatial demand patterns are not mode specific, they illustrate how ride-sourcing vehicles operating that demand are likely to get accumulated within those area if no rebalancing strategy is implemented. Other regions of Lyon, such as the southern industrial neighborhoods may also suffer from a lack of outgoing demand to compensate the ingoing flows.

Note that these incoming flows are fed by flows from the periphery of the area through network entry points not shown on this map.

### 5.1.3 Network specification

The Lyon area overall network is quite large. In order to conduct rapid simulations, we model the traffic on a simplified network of the city. The tertiary roads are eliminated from the network to keep only the primary and secondary urban roads and highways, as illustrated by Figure 9. This simplified network corresponds to a network composed of 5586 links and 3605 nodes, for a total link length of 542 km.



*Figure 9: Study network. Red links correspond to highways, yellow to primary road section and green to secondary road sections.*

### 5.1.4 Request generation

We simulate 4 hours of road traffic, representing the 4 morning peak hours of a typical day, between 6am and 10am. This demand is constructed from a dynamic origin-destination matrix covering the city of Lyon and its urban area for a typical day. The demand is characterized hour by hour, with a spatial granularity defined at the scale of an IRIS zone, the reference spatial division of the French statistical institute. To obtain the flows at the restricted scale of the city, the flows coming from or going to the Lyon periphery are reassigned to entry and exit nodes at the periphery of the network. The application of a Poisson process allows the assignment of departure times to travelers at microsecond resolution, while a uniform drawing of origins and destinations among the nodes of the corresponding IRIS area allows for a finer spatial distribution of the demand.

At this stage, we do not have data on mobility on demand services in the city of Lyon. To simulate the operations of this type of service, we assign a share of the personal vehicle demand to a single mobility on demand service. In this study, this share of demand is arbitrarily set at 15% of internal flows (11% of overall demand). At this stage, ride-sourcing is not considered accessible to flows coming from or going to the outskirts of the city in order to avoid entry and exit of MoD vehicles from the simulated perimeter and to work with a constant fleet.

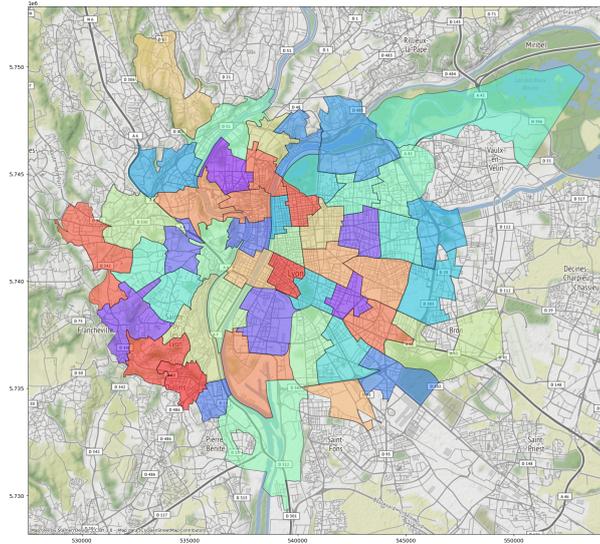


Figure 10: Partitioning of Lyon into 50 demand zones

### 5.1.5 Network partitioning

The methods we implement require the definition of service areas or local control perimeters. These service areas are defined from the clustering of IRIS zone centroids, using the K-means method. We select a partitioning in 50 zones, associated on average to 72 nodes of the network. This partitioning is illustrated in Figure 10.

### 5.1.6 Simulation settings

The numerical experiments presented in this deliverable have been conducted on the simulators developed by EPFL and Univ. Gustave Eiffel. These simulators rely on the same modeling framework, relying on macroscopic traffic model, the trip-based Macroscopic Fundamental Diagram (MFD). For details about this framework, we refer the reader to D6.3, in which we extensively described the overall modeling chain and how it is applied to fleet management.

As the used traffic flow model directly relates the average traffic speed in one area with the number of vehicles accumulated there, ensuring the comparability of the results only takes:

- uniformly designing the individual vehicles and the mobility service strategies (routing, matching);
- defining a common set of parameters;
- setting a unique demand scenario.

Both simulators were set up according to this objective and we checked on a test case without mobility services (individual mobility only) that the traffic dynamics were the same.

In the case of our mobility service implementation, the comparative analysis of results presented below allowed us to identify a technical difference in the design of the matching algorithm, which has an impact on some of the performance of the service. It does not prevent a comprehensive comparison of the proposed algorithms, but will be corrected shortly as part of a continuous improvement of our results.

Table 1 below summarizes the selected simulation parameters.

*Table 1: Selected parameters*

Parameter		Unit	Values
Simulation length		hr	4
Simulation warm-up time <sup>1</sup>		hr	0.5
Ride-sourcing fleet size	$M$	vehicles	4000
Max. matching time <sup>2</sup>	$\Phi$	minutes	1
Rebalancing frequency	$\Delta_{rebalancing}$	minutes	10
Max. waiting time for solo trip	$\Delta_{solo}$	minutes	10
Max. waiting time for pool trip	$\Delta_{pool}$	minutes	20
Max. trip detour ratio <sup>3</sup>	$\Omega$	-	0.2
Max. pool total detour ratio <sup>4</sup>	$\Psi$	-	1
Value of time: constant for all ride-sourcing passengers	$\nu$	\$/hr	25
Solo trip base fare	$f_{s,base}$	\$	2.20
Solo trip per-km fare	$f_{s,v}$	\$/km	1.00
Pool trip base fare	$f_{p,base}$	\$	2.00
Pool trip per-km fare	$f_{p,v}$	\$/km	0.80

In this deliverable, a single fleet size setting was explored, set to 4000 vehicles, yet the impact of fleet size impacts has been studied in previous works [53], [47]. It corresponds to a large fleet size scenario which ensures that some vehicles are available for rebalancing and pickup.

<sup>1</sup>Simulation time after which KPIs will be collected

<sup>2</sup>Time allowed for the ride-sourcing platform to find matching vehicles for a new request, after which the passenger will exit the platform and travel by private vehicle

<sup>3</sup>Ratio between distances of pool and solo trip for one passenger

<sup>4</sup>Ratio between the total distance of a pool trip and the total lengths of respective individual solo trips

## 5.2 Comparative results

### 5.2.1 Selected KPIs

We compare the three strategies proposed through a set of 13 KPIs split into 4 focus categories:

- The Service category looks at the results with a macroscopic scope. The number of passengers served is a key evaluation criteria. We additionally look at the amount of pool rides for scenarios involving ride-splitting. Finally, the dynamic evolution of vehicles states provides indicators regarding vehicle utilization, i.e., whether the fleet is used to its full capacity or not.
- The Users category evaluate the service from the user perspective. We look into waiting and service time dynamics throughout the simulation, as well as overall travel time.
- The Vehicles category evaluate the service from the drivers perspective. We assess the traveled distances splitted into empty and busy distances.
- The final indicator category related to Traffic variable. In each scenario, we evaluate the vehicle accumulation dynamics, and measure the total traveled distances.

These KPIs are summarized in Table 2 below. In the following result section, they are evaluated after a warm-up period of half an hour of simulation. This period is considered sufficient given the short average travel times (between 5 and 10 minutes) observed between 6am and 7am, and the selected rebalancing frequency of 10 minutes.

Given that the methods studied involve little macroscopic stochasticity (passenger-vehicle matching and vehicle routing decisions are deterministic), the results presented in here were obtained on a single simulation? At a microscopic scale, randomness may affect the precise time of occurrence of the request and the mode choices. But considering the spatial and temporal scales that govern our fleet rebalancing strategies, we can expect it has a marginal impact on the results. The robustness of the results will be assessed in continuation of this work.

	Unit	Dyn. vs. Global
<b>Service quality</b>		
Total served passengers	# passengers	global
Pooled rides	users or %	global
No. of vehicles in each state <sup>5</sup>	# vehs	dynamic
<b>Travelers</b>		
Waiting time	s	dynamic
In-vehicle time	s	dynamic
Total travel time	s	dynamic
<b>Vehicles</b>		
Traveled distance	km	global
Service distance <sup>6</sup>	km or %	global
Empty distance	km or %	global
<b>Traffic</b>		
Vehicle accumulation	# vehs	dynamic
Total traveled distances	km	global

Table 2: Selected KPIs

## 5.2.2 Algorithms comparison

In the following section, we refer to the methods discussed above as:

- 1 - RH: Auctioning-based approach for ride-hailing fleet management
- 2 - RH: Hierarchical rebalancing for ride-hailing fleet management
- 2 - RS: Hierarchical rebalancing for ride-splitting fleet management
- 3 - RS: Incentives rebalancing strategy for ride-splitting fleet management

Additionally, the methods 2 and 3 are explored according to different control settings. Regarding method 2, we explore the impact of the non-distributed local repositioning only (2.1), versus the full hierarchical rebalancing process (2.2). For method 3, the performance of the proposed rebalancing criteria is considered without (3.1) and with a pricing strategy (3.2).

### 5.2.2.1 Service

Table 3 reports the performance of each strategy regarding the number of passengers served. Because the implementation of base cases slightly differ from one platform to the other, each strategy is compared against its corresponding base scenario.

<sup>5</sup>Serving, relocating or idle

<sup>6</sup>including pick-up distance

		Passengers served	%
1 - RH	Base	17371.0	
	Rebalancing	17653.0	1.62
2 - RH	Base	19478.0	
	Local repos.	20186.0	3.63
	Full rebalancing	20349.0	4.47
2 - RS	Base	19524.0	
	Local repos.	20208.0	3.50
	Full rebalancing	20911.0	7.10
3 - RS	Base	19990.0	
	Rebalancing	19773.0	-1.09
	Rebalancing and pricing	19985.0	-0.03

*Table 3: Service performances in number of passenger served*

We observed a sometimes significant difference between the performance of a simple service (base case, ride-hailing, no rebalancing) from one simulation platform to another. In particular, we observed that the matching method implemented in the MnMS platform (for method 1 - RH) provided lower service performances than the one implemented in the method 2 - RH study. These undesirable effects are induced by technical choices of development on which we will come back to later. The identification of this problem will allow us to correct it within the project time-frame, for an improvement and a subsequent valorization of these results.

The overall performance comparison also shows uneven results at this stage between the different rebalancing methods, with a gain in passengers served ranging from approximately -1 to +7%. In a ride-hailing setting, the rebalancing strategy based on auctioning is shown to marginally increase the number of passengers served (+282 travelers) compared to the base scenario.

In contrast, the hierarchical rebalancing strategy is shown to increase the number of served users by 871 passengers (+4.47%). In that operational setting, our results illustrate that the local repositioning strategy alone is already able to improve the service performances by 3.63%. This result is confirmed in a ride-splitting operational framework (2 - RS). The local repositioning allows a 3.5% increase of the number of passengers served, while the complete hierarchical strategy allows to reach an increase of 7.10% users served (+1387 users). The benefits of ride-splitting on the service performances compared to ride-hailing is also shown to be larger when applying the hierarchical rebalancing strategy than when no rebalancing policy is implemented.

The last strategy implemented (M3) in a ride-splitting framework shows limited results at this stage. The effects of simple rebalancing are presently counterproductive, with a reduction in the number of passengers served of 1%. This result is discussed in Section 5.2.3.3. However, the pricing strategy is shown to positively compensate for this undesirable effect, and further work on adapting the method to this case study should make it possible to take advantage of the benefits of the

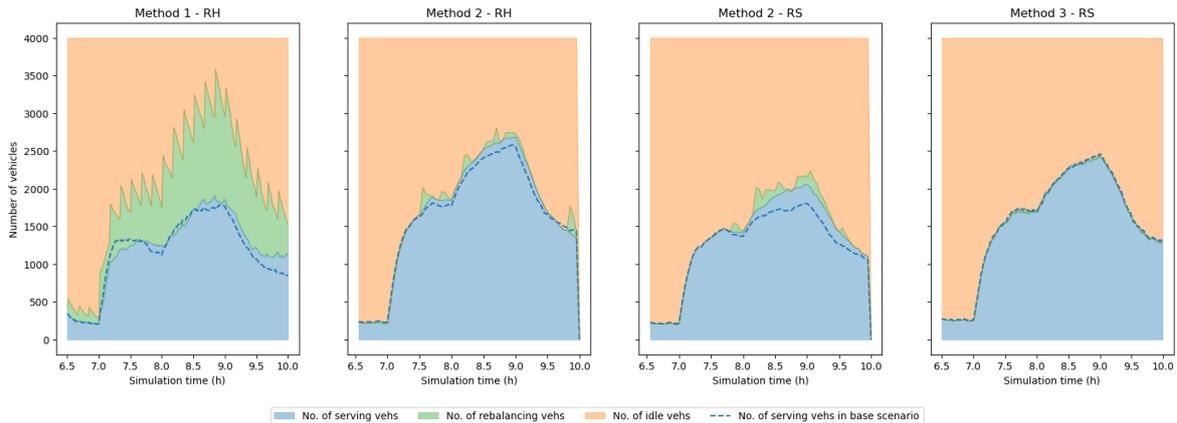


Figure 11: evolution of the number of vehicle in each state (serving, idle or rebalancing) in base and rebalancing scenarios.

method already demonstrated for car-sharing systems [19].

Figure 11 illustrates the evolution of the vehicles states (serving, rebalancing or idle) throughout the four hours of simulation for each rebalancing strategy. For the sake of readability, only the results of the full control strategies are displayed, together with the number of serving vehicles in the base scenario (dashed line).

Regarding Method M1, the limited number of serving vehicles in the base scenario confirms the impact of the matching method discussed above. The study of the rebalancing fleet (green area) also shows that although a large number of vehicles are mobilized by the rebalancing every 10 minutes, it does not increase the number of passengers served in the same proportion. Again, the matching algorithm may be limiting, although there may be other reasons for this result. We discuss this in the discussion section.

Regarding Method M2, the lower number of repositioning vehicles highlights the distributed control strategy. It allows giving instructions to multiple vehicles at a time, which each one will have very small distances to travel in order to fulfil the orders. Therefore, the majority of the vehicles in repositioning activities come from the upper layers, which are responsible for longer distances dispatches. In conclusion, although most vehicles perform quick small rebalancing movements, as ordered in the distributed control, a few vehicles have longer distances to travel in order to comply with near-future demand requests.

Still regarding Method M2, but observing the differences between ride-hailing only and ride-splitting, it shows that fewer vehicles were used to serve requests. This becomes clearer in the peak-hour, where nearly 500 more vehicles were idle. However, it served fewer requests too. With the increased number of rebalancing vehicles, it shows that the rebalancing strategy was trying to compensate the smaller number of vehicles moving (with passengers) to areas where the predictions say they will be needed.

Regarding Method 3, we observe that the rebalancing strategy slightly decreases the number of in-service vehicles, and increases slightly the number of idle vehicles. However, the number of rebalancing vehicles remains low, which is likely due to the formulation and definition of the rebalancing trip destination or origins, resulting in very few vehicles meeting these criteria.

#### 5.2.2.2 Users

In Figure 12 are illustrated the average riding and total travel time for the users of the mobility service during the simulation. Table 4 details the average service, waiting and overall travel time for each strategy. The order of magnitude of these variables are globally similar between test cases, with service time varying between 11.26 (M1 - base) and 12.71 minutes (M2 - RH, base), and waiting time varying between 2.24 (M2 - RS, full rebalancing) and 3.70 (M2 - RH, base).

Taking rebalancing strategies individually, Method 1 is shown to increase both the service time (+1.71%) and the waiting time (+48.55%) of passengers. Overall, the average travel time of passenger increases by +9.75%. This significant increase in travel time is a priori a weakness of the method. However, considering the increased of passenger served, these results suggest a more equitable service of the different service areas, with some users having to wait for pick-up longer so that more users can be served.

In both ride-hailing and ride-splitting settings, Method 2 is shown to reduce the service and waiting time of passengers. The full rebalancing strategy has the strongest impact on the results compared to the corresponding base case scenario. In the ride-splitting operation mode, it allows to reduce the average waiting time by up to 8.93%. Interestingly, in a ride-hailing setting, the local-repositioning-only strategy provides better waiting time reduction than the full rebalancing strategy. However, it should be seen as a consequence of the objectives in each layer of the method. In the one hand, the upper-layers try to minimize the number of lost requests (and minimize the time to fulfil this objective). In the other hand, the lower-layer tries to distribute the vehicles, such that the expected waiting time is minimized. Therefore, the full strategy lowers losses and keep good waiting times, whereas the local strategy only minimizes the waiting times (indirectly serving more requests too). Overall, M2 reduces the average travel time of served users by 1.46% for the ride-hailing service, and by 2.32% for the ride-splitting service.

Regarding Method 3, results in Table 4 show that the rebalancing strategies slightly increase the travel times (0.48% for rebalancing, and 0.80% for rebalancing and pricing), but the waiting times are slightly reduced. This finding can be explained by the difference in the selection of the pooling option. More specifically, the base scenario selects the pooling option with the lowest ratio between total distances of pooled and solo trips, whereas the rebalancing scenarios examine five options with the lowest ratio, and checks in ascending order whether the trip can serve as rebalancing trip. As a result, the suggested options in the base case have lower

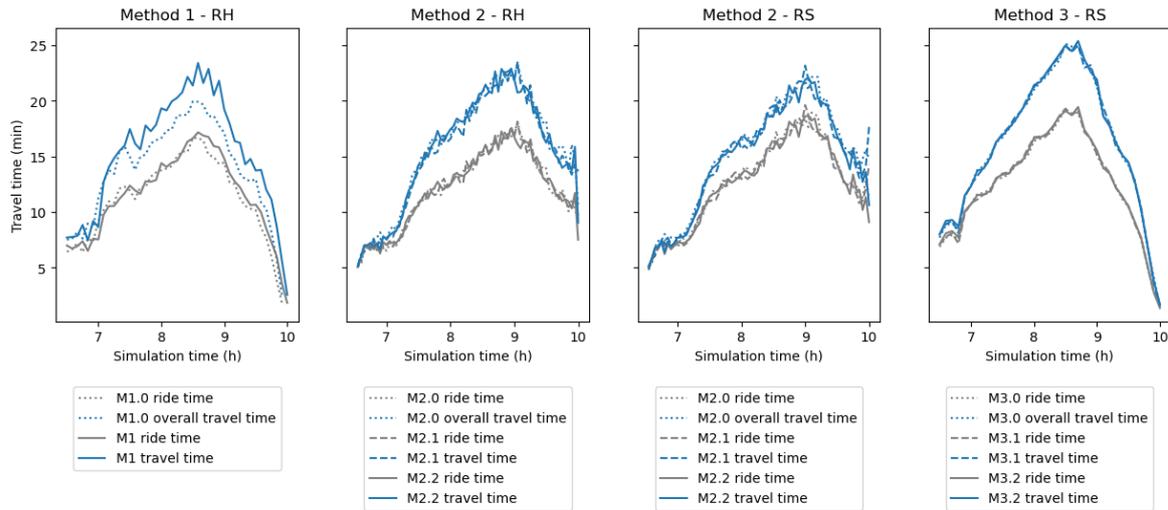


Figure 12: Evolution of average in-vehicle and total travel time in base and rebalancing scenarios.

		Travel time	%	Waiting time	%	Total	%
1 - RH	Base	11.26		2.33		13.59	
	Rebalancing	11.45	1.71	3.46	48.55	14.92	9.75
2 - RH	Base	12.00		3.70		15.70	
	Local repos.	11.99	-0.16	3.36	-8.94	15.35	-2.23
	Full rebalancing	11.97	-0.25	3.50	-5.37	15.47	-1.46
2 - RS	Base	12.71		2.46		15.17	
	Local repos.	12.66	-0.42	2.26	-8.16	14.91	-1.67
	Full rebalancing	12.58	-1.05	2.24	-8.93	14.82	-2.32
3 - RS	Base	12.46		3.66		16.12	
	Rebalancing	12.52	0.48	3.63	-0.87	16.15	0.17
	Rebalancing and pricing	12.56	0.80	3.65	-0.27	16.21	0.56

Table 4: Evolution of average service, waiting and overall travel time for the service users

travel times compared to the rebalancing options. On the other hand, the waiting time reduction in the rebalancing scenarios, although not a direct outcome of the matching algorithm, can be explained by the decrease in vehicle accumulation, which allows vehicles to reach passengers faster than in the base scenario.

### 5.2.2.3 Vehicles

In Table 4 are reported the average distances traveled by vehicles when in service or idle for each of our rebalancing scenarios. The evaluation of these distances is also a good indicator of the performances of the service as they can be used as a proxy for income estimation and service externalities.

Due the matching algorithm designed for studying M1, the limited number of passengers served explains why the base scenario displays shorter average service distance than for the other scenarios. This rebalancing strategy implies an increase

		Avg. distance serving (km)	Avg. distance serving (%)	Avg. distance empty (km)	Avg. distance empty (%)	Avg. total dist. (km)	Avg. total dist. (%)
1 - RH	Base	19.84		0		19.84	
	Rebalancing	20.49	3.25	13.60	+ inf	34.08	71.77
2 - RH	Base	26.86		64.65		91.51	
	Local repos.	27.12	0.98	64.67	0.04	91.80	0.32
2 - RS	Full rebalancing	27.62	2.84	64.23	-0.65	91.85	0.38
	Base	20.90		70.61		91.51	
3 - RS	Local repos.	21.54	3.07	70.28	-0.48	91.82	0.33
	Full rebalancing	22.46	7.47	69.69	-1.30	92.15	0.70
3 - RS	Base	25.05		49.18		74.23	
	Rebalancing	24.83	-0.89	49.30	0.24	74.13	-0.14
	Rebalancing and pricing	25.03	-0.09	49.22	0.07	74.24	0.01

*Table 5: Comparison of vehicles average traveled distances*

of 3.25% of the service distance. Regarding the empty travel distance computation when applying this rebalancing strategy, the comparison with the others scenarios is also limited by the fact that in scenario 1 idle vehicles are considered stopped while the implementation of other scenarios consider cruising idle vehicles. Therefore, when no rebalancing is applied, idle vehicles wait for the next matching at a fixed location and the empty distances are null.

Method 2 allows increase the average in-service distance of vehicles by 2.84% for a ride-hailing service and 7.47% for the ride-splitting service. Simultaneously, the average idle distance reduces slightly for ride-hailing and ride-splitting services. The overall average travel distance rise is limited compared to the improvement of passengers served (7.1%). That is reached by making vehicles closer to passengers, while trying to travel the least without passengers. In one hand, the local strategy only gives short distances for empty vehicles. In the other hand, when vehicles must travel empty for longer distances, the Middle-Layer chooses these vehicles such that the total distance need is shortened (locally and inter-regionally).

We observe that method 3 performs less well. While the distances traveled are relatively unaffected, the method tends to contribute to a slight decrease in the service distance, and a slight increase in the idle distance. The rebalancing strategy alone seems responsible for this effect, while the addition of the pricing strategy balances out the effects of rebalancing and pulls the results in the right direction. These results are discussed further below.

		Max. acc. (# vehs)	%	Total distance (km)	%
1 - RH	Base	34239.0		1060060.87	
	Rebalancing	34238.0	-0.00	1118060.68	5.47
2 - RH	Base	36010.0		1351851.52	
	Local repos.	35674.0	-0.93	1350374.26	-0.11
	Full rebalancing	35713.0	-0.82	1349879.32	-0.15
2 - RS	Base	35956.0		1351885.19	
	Local repos.	35634.0	-0.90	1350417.46	-0.11
	Full rebalancing	35095.0	-2.39	1348834.63	-0.23
3 - RS	Base	37883.0		1244017.88	
	Rebalancing	37840.0	-0.11	1244468.30	0.04
	Rebalancing and pricing	37840.0	-0.11	1243990.45	-0.00

*Table 6: Traffic variables evolution*

#### 5.2.2.4 Traffic

Finally, we study if the strategies implemented have an impact on the overall road traffic. Table 6 displays the maximal vehicle accumulation and total distances traveled in each scenario. The results account for both the personal vehicles and the ride-sourcing fleet. At this stage, the impact on the traffic variables is shown to be very limited. This can be explained by the high share (89%) of demand for individual mobility (versus 11% ride-sourcing requests) in the demand scenarios studied here, that contributes to an unavoidable traffic. By studying different demand scenarios, our future work will further evaluate the potential of these services to reduce road traffic and its externalities.

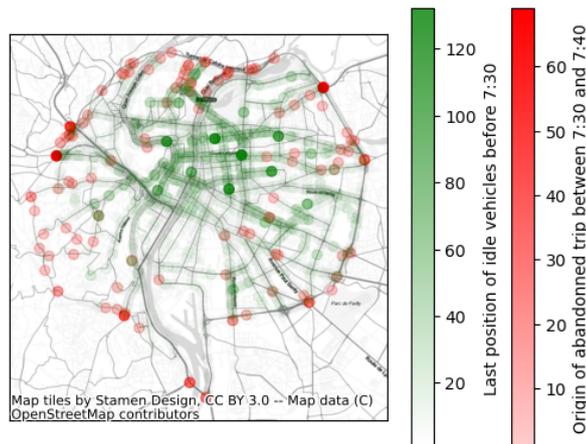


Figure 13: Analysis of abandoned trips between 7:30 and 7:40am

## 5.2.3 Discussion

### 5.2.3.1 Method 1

Several of the results of the M1 method presented can be considered as unsatisfactory.

First, the number of passengers served in the case of a baseline scenario (without rebalancing) is very limited, and lower than the results achieved by the baseline scenarios of methods 2 and 3. This difference allowed us to identify a limitation of the matching algorithm implemented in the MnMS platform. This is a technical design choice whereby the strategy of matching passengers with the nearest vehicle is based on the Euclidean distance calculation rather than a network distance calculation. While this approach allows fast computations and satisfactory results on densely connected networks, it was found to be deficient here. We identified that the nearest vehicle based on the Euclidean distance sometimes had to travel long distances to reach the passenger (ring road, bridge crossing, ...). This implied long waiting times for passengers who finally decided to abandon the ride and take their personal vehicle. Our current explorations show that this feature that affects the baseline scenario also impacts the rebalancing scenario. Thus, the relocated vehicles are not necessarily assigned to the new requests, limiting the gain in passengers served.

Another issue is the large number of vehicles relocating, and the impact of this on empty distances traveled. It seems that depending on the prediction horizon chosen, the utility function used in these simulations may favor the rebalancing of distant vehicles, contributing to a large number of empty vehicles. This will be the subject of further study, in which a prioritisation of short rebalancing distances by local controllers may compensate this problem.

### 5.2.3.2 Method 2

Although highly optimized, the proposed multi-layer control framework still resulted in a number of unserved requests, while there were drivers idling. Figure 14 highlights the visualization of ride-sourcing requests and focus on the difference between the served and unserved requests. One can note that, at one point, the central area of Lyon is well-served and faced no losses. On the other hand, requests coming from outside the city center, especially from the western and southern parts of Lyon, could not be served.

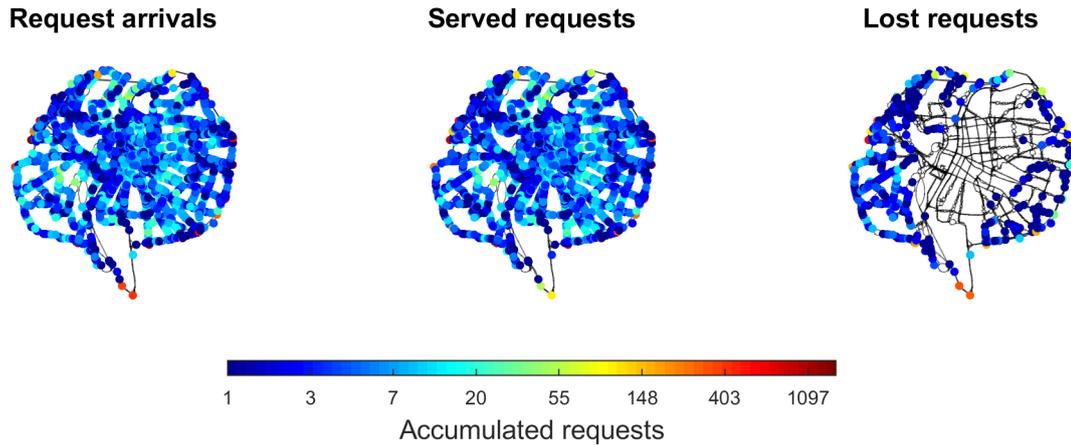


Figure 14: Geographical summary of ride-sourcing requests service quality using the Multi-Layer Control.

To understand what caused this problem, Figure 15 details the unserved requests to the westernmost macroscopic region used for the proposed Multi-Layer Control approach. Firstly, it highlights the imbalance between the demand and the supply of drivers in that area. While the lost requests are concentrated on the westernmost part of Region 1, idle drivers were concentrated on the easternmost part of Region 1. The design of the regions used a weighted-network  $k$ -means algorithm with three random initial centroids. We have to highlight a few points about this result:

- The outcome of this regional design made the westernmost region accessible only by a small set of edges (reaching 23 intersections, out of the more than 700 in that area).
- Most of the demand destinations (area where drivers tend to become available for new requests) concentrate in the central area, especially close to 5 of the 23 intersections connecting to Regions 2 and 3.

Recall that the Upper- and Middle-Layers of the approach dispatch the selected vehicles to the border of the destination area, before the Lower-Layer acts with its distributed coverage control. Therefore, many vehicles accumulate near the border because of the few entrances and proximity to high demand area. However, coverage control cannot differentiate vehicles at the same location, since every intersection is equally distant to overlapping centroids. To work around this issue, we

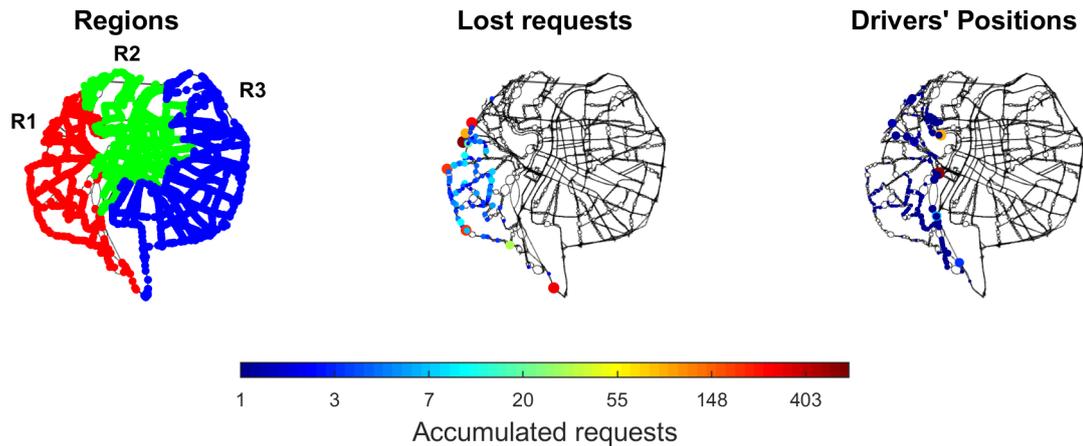


Figure 15: (Left) Macroscopic regional setting used in the Multi-Layer Control approach. (Center) Concentration of lost requests in Region 1 (R1). (Right) Average concentration of idle ride-sourcing vehicles in Region 1 (R1).

adapted the coverage control to ignore overlapping positions and consider only one of them in the process. Therefore, a single vehicle is gets instructions at a time out of the overlapping ones. However, the dispatch of a single driver was not enough to serve the demand in time, resulting in drivers being matched long before reaching the repositioning destination. In the meantime, the other drivers wait for the next control step to get new instructions, and new drivers enter the area, increasing the waiting line of idle vehicles even more.

There are possible alternatives to overcome this limitation with the coverage control. A indirect solution is to rethink the regional design. One alternative is the development of methods that account for the accessibility of regions (providing enough connecting edges to neighbouring areas). Previous works in MFD clustering could be utilized in this direction [54], [55]. Another alternative is to rethink the coverage control layer to dispatch multiple overlapping vehicles at once, instead of the current one at a time policy. Such alternative requires attention not to transfer the entire supply of drivers at once, and transferring the imbalance to another area.



Figure 16: Demand and loss aggregated at the centroids of 50 zones from pool scenario of Method 3

### 5.2.3.3 Method 3

The proposed rule-based rebalancing strategy shows limited improvement at this stage. Figure 16 illustrates the demand distribution aggregated over the 50 region centroids, where the marker size corresponds to the demand intensity. It can be observed that while the peripheral regions are characterized by higher demand outflow indicated by the radius of blue circles, the central region experiences higher vehicle inflow illustrated by the sizes of the yellow circles. In the same figure, the size of blue crosses indicates the aggregated number of abandoned requests from the base case scenario for pooling after the four-hour simulation run. In fact, the rebalancing scenario, although targets areas of high trip abandonment, still showed higher demand losses in peripheral areas, as in Figure 16.

The proposed ride-splitting rebalancing strategies described in Section 4.3 currently shows limited improvement, as compared earlier in Table 3. Additionally, Figure 11 indicates that the number of rebalancing vehicles at any time is very low (there are 25 rebalancing vehicles at the maximum).

One immediate consideration for service improvement is to implement an additional rebalancing objective at trip origins in low demand areas. So far, the rebalancing strategy only targets trip destinations, aiming to extend the final destination of solo trips towards high demand areas. For the next step, we intend to incorporate a decision rule to discourage pooling if two trips originate in low demand areas and have destinations in high demand areas. The hypothesis here is that by aiming to dispatch as many empty vehicles from low demand regions to high demand region through serving solo trips whenever possible, such vehicles can be employed in a similar manner compared to the empty vehicle relocation strategy in Method 2, but with a reduction in distances travelled while empty; in turn, we

can investigate whether this can increase the vehicle supply at higher demand regions in peripheral areas in the network.

## 6 Conclusion

The asymmetric commuting mobility patterns observed at the urban scale affect the fleets of on-demand mobility services. During the morning peak, certain zones attract massively passenger flows: university, multimodal hubs, industrial and employment areas (see Section 5.1.2). Without applying fleet reorganization strategies, on-demand service vehicles, subject to these demand patterns, may see their fleet accumulate in these areas and experience a shortage of vehicles in high-demand regions instead.

Anticipating these patterns and proactively reorganizing fleets is a guarantee for minimizing passenger waiting times, satisfying a broader demand, and improving service efficiency. In this deliverable, we first discuss the issue of demand prediction for these mobility services and present a work conducted as part of the DIT4TraM project. This research supports the methodological development of proactive rebalancing strategies.

We then propose three algorithms for rebalancing a fleet of on-demand mobility services:

- A first algorithm (M1) is based on a game theoretical approach via a bidding and matching process. It is studied at this stage in the context of ride-hailing-only operations.
- A second algorithm (M2) relies on a hierarchical rebalancing approach based on an upper-level model predictive control framework and a local fleet repositioning. This algorithm, adapted to ride-hailing and ride-splitting operations, is studied in both contexts.
- A third algorithm (M3) leverages ride splitting to limit vehicle accumulation in low-demand areas and foster vehicle in-service repositioning in high-demand areas. This algorithm, based in essence on ride-splitting, is only studied in this context.

All three rebalancing algorithms presented in this deliverable were developed within the DIT4TraM project. For the first time, these methods are tested at a large scale and compared in a common case study. The city of Lyon, France, is picked as a study area and used to simulate 4 hours of the morning peak traffic.

The comparative analysis we conduct focuses on criteria of overall service performance, passenger travel times, vehicle distances traveled, and conventional traffic variables. They allow us to underline gains in terms of number of passengers served and reduction of waiting times or empty kilometers traveled. Yet, the results obtained on other case studies suggest that these first results obtained on the city of Lyon do not allow us to illustrate the full potential of the performances of the developed methods [56], [53], [47], [57]. The previous part of this document proposed explanatory elements for the limitations encountered at this stage. They can be summarized as follows:

1. The algorithmic design of some platform modules;
2. Methodological choices specific to each method (utility functions, coverage control, pricing strategies);
3. The network topology and the demand patterns characterized by a high inflow from the outskirts of the network;
4. Some calibration elements, such as zonal partitioning.

Continuing the efforts of calibration the methods to the case study presented here, conducting extended sensitivity analyses, and exploring the proposed research directions should allow an improvement of these results. Both teams (EPFL and Univ. Gustave Eiffel) will continue to work together within the project to improve the algorithms and look for cross-fertilization. In particular, the joint use of these different methods will be explored, for instance to more systematically use microscopic fleet repositioning or pricing strategies.

## Publications for D3.3

- [1] C.V. Beojone, P. Zhu, I. Sirmatel, and N. Geroliminis, “A Hierarchical Control Framework for Vehicle Repositioning in Ride-hailing systems,” in *25th International Symposium on Transportation and Traffic Theory (ISTTT 25)*, 2024 (under review).
- [2] M. Khalesian, A. Furno, and L. Leclercq, “High-Resolution Traffic Demand Prediction, Based on Deep Learning and Hierarchical Reconciliation,” *SSRN Electronic Journal*, 2022 (under review). DOI: [10.2139/ssrn.4250272](https://doi.org/10.2139/ssrn.4250272).
- [3] C.V. Beojone and N. Geroliminis, “Providing a Revenue-forecasting Information Scheme as an Incentive to Relocate Compliant Ride-Sourcing Drivers,” in *102nd Transportation Research Board Annual Meeting (TRB 2023)*, 2023.
- [4] M. Seppecher and L. Leclercq, “A Decentralised Auction-Based Framework for Rebalancing Ride-Hailing Fleets with Uncertain Request Probabilities,” in *102nd Transportation Research Board Annual Meeting (TRB 2023)*, Washington D.C., USA, Jan. 2023.
- [5] C. V. Beojone and N. Geroliminis, “A dynamic multi-region MFD model for ride-sourcing systems with ridesplitting,” *arXiv*, vol. 2211.14560, 2022. DOI: [10.48550/arXiv.2211.14560](https://doi.org/10.48550/arXiv.2211.14560).
- [6] C.V. Beojone and N. Geroliminis, “A dynamic multi-region MFD model for ride-sourcing systems with ridesplitting,” in *101st Transportation Research Board Annual Meeting (TRB 2022)*, Washington D.C., USA, 2022.
- [7] C.V. Beojone and N. Geroliminis, “An optimized driver repositioning strategy in ridesplitting with earning estimates: A two-layer dynamic model and control,” in *10th Symposium of the European Association for Research in Transportation (hEART 2022)*, Leuven, Belgium, 2022.
- [8] C.V. Beojone and N. Geroliminis, “Repositioning ridesplitting vehicles through pricing: A two-region simulated study,” in *11th Triennial Symposium on Transportation Analysis Conference (TRISTAN XI)*, Mauritius, 2022.
- [9] M. Wang and N. Geroliminis, “Ride-sourcing fleet rebalancing with proactive and targeted pooling incentives and surge pricing,” in *22nd Swiss Transport Research Conference (STRC 2022)*, Ascona, Switzerland, 2022.
- [10] M. Wang and N. Geroliminis, “The impact of proactive ride-splitting incentives on revenue and service quality,” in *10th Symposium of the European Association for Research in Transportation (hEART 2022)*, Leuven, Belgium, 2022.

## References

- [1] A. Wallar, M. Van Der Zee, J. Alonso-Mora, and D. Rus, "Vehicle Rebalancing for Mobility-on-Demand Systems with Ride-Sharing," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, ISSN: 2153-0866, Oct. 2018, pp. 4539–4546. DOI: [10.1109/IRoS.2018.8593743](https://doi.org/10.1109/IRoS.2018.8593743).
- [2] A. Simonetto, J. Monteil, and C. Gambella, "Real-time city-scale ridesharing via linear assignment problems," *Transportation Research Part C: Emerging Technologies*, vol. 101, pp. 208–232, 2019. DOI: [10.1016/j.trc.2019.01.019](https://doi.org/10.1016/j.trc.2019.01.019).
- [3] Y. Liu and S. Samaranayake, "Proactive rebalancing and speed-up techniques for on-demand high capacity ridesourcing services," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–8, 2020. DOI: [10.1109/TITS.2020.3016128](https://doi.org/10.1109/TITS.2020.3016128).
- [4] F. Miao, S. Han, S. Lin, *et al.*, "Taxi Dispatch with Real-Time Sensing Data in Metropolitan Areas: A Receding Horizon Control Approach," *Proceedings of the ACM/IEEE Sixth International Conference on Cyber-Physical Systems*, pp. 100–109, Apr. 2015. DOI: [10.1145/2735960.2735961](https://doi.org/10.1145/2735960.2735961). arXiv: [1603.04418](https://arxiv.org/abs/1603.04418).
- [5] F. Miao, S. Han, S. Lin, *et al.*, "Data-Driven Robust Taxi Dispatch under Demand Uncertainties," *arXiv:1603.06263 [cs]*, Oct. 2017. arXiv: [1603.06263 \[cs\]](https://arxiv.org/abs/1603.06263).
- [6] J. Alonso-Mora, A. Wallar, and D. Rus, "Predictive routing for autonomous mobility-on-demand systems with ride-sharing," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017, pp. 3583–3590. DOI: [10.1109/IRoS.2017.8206203](https://doi.org/10.1109/IRoS.2017.8206203).
- [7] P. Zhu, I. I. Sirmatel, G. Trecate, and N. Geroliminis, "Distributed coverage control for vehicle rebalancing in mobility-on-demand systems," *TRB Annual Meeting*, pp. 22–03 340, 2022.
- [8] M. Ramezani and M. Nourinejad, "Dynamic modeling and control of taxi services in large-scale urban networks: A macroscopic approach," *Transportation Research Part C*, vol. 94, pp. 203–219, 2018. DOI: [10.1016/j.trc.2017.08.011](https://doi.org/10.1016/j.trc.2017.08.011).
- [9] Y. H. Wu, L. J. Guan, and S. Winter, "Peer-to-Peer Shared Ride Systems," in *GeoSensor Networks: Second International Conference, GSN 2006, Boston, MA, USA, October 1–3, 2006, Revised Selected and Invited Papers*, ser. Lecture Notes in Computer Science, S. Nittel, A. Labrinidis, and A. Stefanidis, Eds., Berlin, Heidelberg: Springer, 2008, pp. 252–270. DOI: [10.1007/978-3-540-79996-2\\_14](https://doi.org/10.1007/978-3-540-79996-2_14).
- [10] A. Manjunath, V. Raychoudhury, S. Saha, S. Kar, and A. Kamath, "CARE-Share: A Cooperative and Adaptive Strategy for Distributed Taxi Ride Sharing," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–17, 2021. DOI: [10.1109/TITS.2021.3066439](https://doi.org/10.1109/TITS.2021.3066439).
- [11] M. Nourinejad and M. J. Roorda, "Agent based model for dynamic ridesharing," *Transportation Research Part C: Emerging Technologies*, vol. 64, pp. 117–132, 2016. DOI: [10.1016/j.trc.2015.07.016](https://doi.org/10.1016/j.trc.2015.07.016).
- [12] L. Chen, A. H. Valadkhani, and M. Ramezani, "Decentralised cooperative cruising of autonomous ride-sourcing fleets," *Transportation Research Part C: Emerg-*

- ing Technologies*, vol. 131, p. 103336, Oct. 2021. DOI: [10.1016/j.trc.2021.103336](https://doi.org/10.1016/j.trc.2021.103336).
- [13] S. M. Meshkani and B. Farooq, “Decentralized Matching in Shared Intelligent Vehicles Fleet,” Feb. 2022. DOI: [10.48550/arXiv.2202.01121](https://doi.org/10.48550/arXiv.2202.01121).
- [14] A. Lu, P. Frazier, and O. Kislev, “Surge pricing moves uber’s driver partners,” *Social Science Research Network*, vol. 3180246, 2018. DOI: [10.2139/ssrn.3180246](https://doi.org/10.2139/ssrn.3180246).
- [15] A. Sadeghi and S. L. Smith, “On re-balancing self-interested agents in ride-sourcing transportation networks,” in *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2019, pp. 5119–5125. DOI: [10.1109/CDC40024.2019.9030043](https://doi.org/10.1109/CDC40024.2019.9030043).
- [16] J. W. Powell, Y. Huang, F. Bastani, and M. Ji, “Towards reducing taxicab cruising time using spatio-temporal profitability maps,” in *Advances in Spatial and Temporal Databases*, D. Pfoser, Y. Tao, K. Mouratidis, et al., Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 242–260.
- [17] J. Castillo, D. Knoepfle, and E. Weyl, “Surge pricing solves the wild goose chase,” *Social Science Research Network*, vol. 2890666, 2018. DOI: [10.2139/ssrn.2890666](https://doi.org/10.2139/ssrn.2890666).
- [18] G. H. D. A. Correia, D. R. Jorge, and D. M. Antunes, “The added value of accounting for users’ flexibility and information on the potential of a station-based one-way car-sharing system: An application in lisbon, portugal,” *Journal of Intelligent Transportation Systems*, vol. 18, no. 3, pp. 299–308, 2014. DOI: [10.1080/15472450.2013.836928](https://doi.org/10.1080/15472450.2013.836928). eprint: <https://doi.org/10.1080/15472450.2013.836928>. [Online]. Available: <https://doi.org/10.1080/15472450.2013.836928>.
- [19] P. Stokkink and N. Geroliminis, “Predictive user-based relocation through incentives in one-way car-sharing systems,” *Transportation Research Part B: Methodological*, vol. 149, pp. 230–249, 2021. DOI: <https://doi.org/10.1016/j.trb.2021.05.008>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0191261521000874>.
- [20] M. Khalesian, A. Furno, and L. Leclercq, “High-Resolution Traffic Demand Prediction, Based on Deep Learning and Hierarchical Reconciliation,” *SSRN Electronic Journal*, 2022 (under review). DOI: [10.2139/ssrn.4250272](https://doi.org/10.2139/ssrn.4250272).
- [21] A. Li and K. W. Axhausen, “Comparison of short-term traffic demand prediction methods for transport services,” *Arbeitsberichte Verkehrs- und Raumplanung*, vol. 1447, 2019.
- [22] J. Xu, R. Rahmatizadeh, L. Bölöni, and D. Turgut, “Real-time prediction of taxi demand using recurrent neural networks,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 8, pp. 2572–2581, 2017.
- [23] H. Luo, J. Cai, K. Zhang, R. Xie, and L. Zheng, “A multi-task deep learning model for short-term taxi demand forecasting considering spatiotemporal dependences,” *Journal of Traffic and Transportation Engineering (English Edition)*, vol. 8, no. 1, pp. 83–94, 2021.
- [24] J. Ke, H. Zheng, H. Yang, and X. M. Chen, “Short-term forecasting of passenger demand under on-demand ride services: A spatio-temporal deep learn-

- ing approach,” *Transportation research part C: Emerging technologies*, vol. 85, pp. 591–608, 2017.
- [25] H. Yang, C. W. Leung, S. C. Wong, and M. G. Bell, “Equilibria of bilateral taxi–customer searching and meeting on networks,” *Transportation Research Part B: Methodological*, vol. 44, no. 8–9, pp. 1067–1083, 2010.
- [26] X. Li, G. Pan, Z. Wu, *et al.*, “Prediction of urban human mobility using large-scale taxi traces and its applications,” *Frontiers of Computer Science*, vol. 6, pp. 111–121, 2012.
- [27] L. Moreira–Matias, J. Gama, M. Ferreira, J. Mendes–Moreira, and L. Damas, “Predicting taxi–passenger demand using streaming data,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 3, pp. 1393–1402, 2013.
- [28] X. Yin, G. Wu, J. Wei, Y. Shen, H. Qi, and B. Yin, “Deep learning on traffic prediction: Methods, analysis, and future directions,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 6, pp. 4927–4943, 2021.
- [29] J. Guan, W. Wang, W. Li, and S. Zhou, “A unified framework for predicting kpis of on-demand transport services,” *IEEE access*, vol. 6, pp. 32 005–32 014, 2018.
- [30] W. Li, J. Cao, J. Guan, *et al.*, “A general framework for unmet demand prediction in on-demand transport services,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 8, pp. 2820–2830, 2018.
- [31] Y. Li, J. Lu, L. Zhang, and Y. Zhao, “Taxi booking mobile app order demand prediction based on short-term traffic forecasting,” *Transportation Research Record*, vol. 2634, no. 1, pp. 57–68, 2017.
- [32] D. Salinas, M. Bohlke–Schneider, L. Callot, R. Medico, and J. Gasthaus, “High-dimensional multivariate forecasting with low-rank gaussian copula processes,” *Advances in neural information processing systems*, vol. 32, 2019.
- [33] N. Davis, G. Raina, and K. Jagannathan, “Grids versus graphs: Partitioning space for improved taxi demand–supply forecasts,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 10, pp. 6526–6535, 2020.
- [34] L. Liu, Z. Qiu, G. Li, Q. Wang, W. Ouyang, and L. Lin, “Contextualized spatial–temporal network for taxi origin–destination demand prediction,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 10, pp. 3875–3887, 2019.
- [35] H. Nguyen, L.–M. Kieu, T. Wen, and C. Cai, “Deep learning methods in transportation domain: A review,” *IET Intelligent Transport Systems*, vol. 12, no. 9, pp. 998–1004, 2018.
- [36] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, “The graph neural network model,” *IEEE transactions on neural networks*, vol. 20, no. 1, pp. 61–80, 2008.
- [37] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *nature*, vol. 323, no. 6088, pp. 533–536, 1986.
- [38] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

- [39] K. Cho, B. Van Merriënboer, D. Bahdanau, and Y. Bengio, “On the properties of neural machine translation: Encoder–decoder approaches,” *arXiv preprint arXiv:1409.1259*, 2014.
- [40] Q. Cheng, Y. Liu, W. Wei, and Z. Liu, “Analysis and forecasting of the day-to-day travel demand variations for large-scale transportation networks: A deep learning approach,” *Transportation Analytics Contest, Tech. Rep*, 2016.
- [41] J. Guo, B. M. Williams, and B. L. Smith, “Data collection time intervals for stochastic short-term traffic flow forecasting,” *Transportation Research Record*, vol. 2024, no. 1, pp. 18–26, 2007.
- [42] Z. Li, Z. Zheng, and S. Washington, “Short-term traffic flow forecasting: A component-wise gradient boosting approach with hierarchical reconciliation,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 12, pp. 5060–5072, 2019.
- [43] C. Oh, S. G. Ritchie, and J.-S. Oh, “Exploring the relationship between data aggregation and predictability to provide better predictive traffic information,” *Transportation Research Record*, vol. 1935, no. 1, pp. 28–36, 2005.
- [44] E. Vlahogianni and M. Karlaftis, “Temporal aggregation in traffic data: Implications for statistical characteristics and model choice,” *Transportation Letters*, vol. 3, no. 1, pp. 37–49, 2011.
- [45] R. Hollyman, F. Petropoulos, and M. E. Tipping, “Understanding forecast reconciliation,” *European Journal of Operational Research*, vol. 294, no. 1, pp. 149–160, 2021.
- [46] P. Mancuso, V. Piccialli, and A. M. Sudoso, “A machine learning approach for forecasting hierarchical time series,” *Expert Systems with Applications*, vol. 182, p. 115 102, 2021.
- [47] C. V. Beojone and N. Geroliminis, “A dynamic multi-region MFD model for ride-sourcing systems with ridesplitting,” *arXiv*, vol. 2211.14560, 2022. DOI: [10.48550/arXiv.2211.14560](https://doi.org/10.48550/arXiv.2211.14560).
- [48] M. L. Fisher and R. Jaikumar, “A generalized assignment heuristic for vehicle routing,” *Networks*, vol. 11, no. 2, pp. 109–124, 1981. DOI: [10.1002/net.3230110205](https://doi.org/10.1002/net.3230110205).
- [49] P. Zhu, I. I. Sirmatel, G. Ferrari-Trecate, and N. Geroliminis, “Idle-vehicle rebalancing coverage control for ride-sourcing systems,” *the European Control Conference 2022 (ECC22)*, 2022.
- [50] J. W. Durham, R. Carli, P. Frasca, and F. Bullo, “Discrete partitioning and coverage control for gossiping robots,” *IEEE Transactions on Robotics*, vol. 28, no. 2, pp. 364–378, 2012.
- [51] C. Beojone and N. Geroliminis, “On the inefficiency of ride-sourcing services towards urban congestion,” *Transportation Research Part C: Emerging Technologies*, vol. 124, p. 102 890, Mar. 2021. DOI: [10.1016/j.trc.2020.102890](https://doi.org/10.1016/j.trc.2020.102890).
- [52] G. Mariotte, L. Leclercq, S. Batista, J. Krug, and M. Paipuri, “Calibration and validation of multi-reservoir MFD models: A case study in Lyon,” *Transportation Research Part B: Methodological*, vol. 136, pp. 62–86, Jun. 2020. DOI: [10.1016/j.trb.2020.03.006](https://doi.org/10.1016/j.trb.2020.03.006).

- [53] M. Seppacher and L. Leclercq, “A Decentralised Auction-Based Framework for Rebalancing Ride-Hailing Fleets with Uncertain Request Probabilities,” in *102nd Transportation Research Board Annual Meeting (TRB 2023)*, Washington D.C., USA, Jan. 2023.
- [54] M. Saeedmanesh and N. Geroliminis, “Clustering of heterogeneous networks with directional flows based on “snake” similarities,” *Transportation Research Part B: Methodological*, vol. 91, pp. 250–269, 2016. DOI: [10.1016/j.trb.2016.05.008](https://doi.org/10.1016/j.trb.2016.05.008).
- [55] E. Hans, N. Chiabaut, and L. Leclercq, “Clustering approach for assessing the travel time variability of arterials,” *Transportation Research Record*, vol. 2422, no. 1, pp. 42–49, 2014. DOI: [10.3141/2422-05](https://doi.org/10.3141/2422-05).
- [56] C.V. Beojone, P. Zhu, I. Sirmatel, and N. Geroliminis, “A Hierarchical Control Framework for Vehicle Repositioning in Ride-hailing systems,” in *25th International Symposium on Transportation and Traffic Theory (ISTTT 25)*, 2024 (under review).
- [57] M. Wang and N. Geroliminis, “The impact of proactive ride-splitting incentives on revenue and service quality.,” in *10th Symposium of the European Association for Research in Transportation (hEART 2022)*, Leuven, Belgium, 2022.



Distributed Intelligence & Technology  
for Traffic & Mobility Management



This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement no. 953783.