A DECENTRALISED AUCTION-BASED FRAMEWORK FOR REBALANCING RIDE-HAILING FLEETS WITH UNCERTAIN REQUEST PROBABILITIES

Manon Seppecher, Ludovic Leclercq

AUCTION-BASED DISTRIBUTED RIDE-HAILING REBALANCING: THE CASE STUDY OF LYON, FRANCE

Manon Seppecher*
Univ. Gustave Eiffel, ENTPE, Univ. Lyon, LICIT-ECO7
UMR_T 9401, F-69518, Lyon, France
Email: manon.seppecher@univ-eiffel.fr

Ludovic Leclercq
Univ. Gustave Eiffel, ENTPE, Univ. Lyon, LICIT-ECO7
UMR_T 9401, F-69518, Lyon, France
Email: ludovic.leclercq@univ-eiffel.fr

* Corresponding author

Paper submitted for presentation at the 103nd Annual Meeting Transportation Research Board, Washington D.C., January 2024

Speciall call by ACP50 committee on "Traffic flow operations in large-scale urban networks"

Word Count: 5688 words + 5 table(s) × 250 = 6938 words

Submission Date: August 1, 2023

ABSTRACT

Ride-sourcing services offer cost-efficient, convenient, and flexible mobility solutions, making them competitive compared to traditional transportation modes. However, these mobility services must address various operational issues to remain attractive. One such challenge arises from fast-changing asymmetric mobility patterns, leading to the accumulation of idle vehicles in attractive zones and shortages of vehicles in other neighborhoods. As a result, passengers experience high pick-up waiting times and abandonment rates. Based on demand anticipation, proactive fleet re-balancing helps counteract the impact of demand patterns on fleet distribution and ensure service continuity. In previous work, we developed a distributed rebalancing method based on negotiation between vehicles and local controllers and analyzed its features on a small abstract network. This paper further investigates the characteristics of the strategy on a larger scale using an actual urban case study, the city of Lyon, France. Extensive sensitivity analyses are conducted, covering various aspects such as user patience, fleet size, uncertainty, and depot locations, among others.

Keywords: fleet management, on-demand mobility, ride-hailing, distributed algorithms, multi-agent modelling, demand uncertainty

INTRODUCTION

Urban environments are witnessing a remarkable diversification of mobility services, with innovative offerings such as AI-powered ride-sourcing and vehicle sharing complementing and competing against conventional public transit systems. These services significantly expand transportation options for urban dwellers and meet an increasingly dynamic and non-regular mobility demand.

In particular, e-hailing services grant urban residents enhanced flexibility and convenience. Compared to public transit, they offer on-spot, on-demand, no-connection rides. Compared to private car use, they are cost-effective and solve parking issues. These individual benefits extend to the collective and environmental scale, as reducing car ownership relieves pressure on land and the need for parking spaces. And while these services do contribute to congestion, mitigating measures such as ride-splitting and limiting empty cruising behavior can address this drawback [1].

However, unlocking these benefits requires urban authorities and service operators to implement effective service management strategies. Public authorities must regulate competition between mobility services, which can lead to increased cruising distances [2]. Meanwhile, service operators must successfully address various operational challenges, including demand prediction, ride pricing, passenger-vehicle matching, vehicle routing, and fleet rebalancing.

Fleet rebalancing consists in providing relocating directions to idle vehicles. In doing so, companies mitigate fleet shortages and prevent service gaps resulting from the large-scale urban mobility dynamics and asymmetric patterns and consequent accumulation of end-of-travel vehicles in attractive areas. A simple rebalancing strategy involves relocating idle vehicles to areas where recent demand was unmet [3–5] have proposed approaches in this direction. However, these strategies rely on past demand information and are reactive. Proactive strategies, based on future demand knowledge or prediction, can achieve better performances. Numerous literature studies have developed proactive rebalancing strategies [6–9]. However, most offer centralized methods, raising questions about scalability, robustness to system failures, and their suitability to model and manage increasingly decentralized services, where self-employed drivers make their own decisions.

In this respect, distributed approaches are attractive alternatives. Recent works have looked into fleet rebalancing through the lens of passengers and drivers-intended incentives, with pricing and information-sharing strategies [10], coverage control [11] or auctioning [12]. In this latter work, we proposed a strategy in which the relocation directions provided to drivers result from a negotiation process with local controllers, which aim at attracting vehicles nearby. While auctioning approaches have previously been employed for developing reactive vehicle-passenger matching strategies [9, 13, 14], this study was the first attempt, to our knowledge, to extend its application to large-scale repositioning. In the present paper, we apply this methodology to a large-scale case study in Lyon, France. We conduct extensive sensitivity analyses to passengers waiting time tolerance, fleet size, demand uncertainty, and depot locations. The results evidence convincing benefits and provide guidelines for further improving the strategy.

This paper is organized as follows. Section 3 summarizes the key elements of the methodology. Section 4 presents the simulation platform and test case. Section 5 details our results. Section 6 concludes this paper.

METHODOLOGY

Outline

We investigate the performance of a decentralized architecture for large-scale, e-hailing fleet rebalancing. In this setup [12], the rebalancing decision-making process is outsourced from ride-sourcing companies' management to a network of controllers that act on behalf of a public authority. The operating zone is partitioned into service areas, each managed by a local controller, supported by physical infrastructures (*e.g.*, taxi stations, vehicle deposits, or parking lots).

Local controllers aim to ensure prompt pick-up in response to local ride-hailing demand. To achieve this, they:

1. Forecast local ride-hailing demand.
2. Issue *relocation offers* accordingly.
3. Engage in negotiating with vacant vehicles to have them relocate within the service area boundaries.

Local demand prediction enables controllers to estimate the number of vehicles required to meet the expected local demand. They periodically (e.g., every 10 minutes) engage in a distributed negotiation process with vehicles to attract the necessary number within their boundaries. This process is modeled as a two-sided matching market and accounts for both the regions and the vehicles' utilities. Once matched with a service area, vehicles relocate to the nearest vehicle depot of this area and wait for a match with a passenger there.

Since a decentralized public authority handles the rebalancing decisions, the method is suited for managing vehicles' rebalancing for multiple and competing ride-sourcing services. However, this paper assumes a monopoly setting with a single company operating on the network. We also assume that vehicles and the ride-sourcing service comply with the rebalancing decisions. Future research may explore the model's properties in a competitive setting and with partially compliant vehicles or services.

The remaining of this section is organized as follows. Section 3.2 presents how the demand prediction allows to anticipate local needs for vacant vehicles and to prepare the negotiation phase. Section 3.3 presents the negotiation process. Section 3.3 defines the regional and vehicle utility applied in this paper.

Relocation offers

Local controllers issue relocation offers as a vector of information and attraction to vacant vehicles. Let i denote a service area and T a rebalancing period. For all k in $\mathbb{N}^*$, we define the k^{th} relocation offer to region i and period $T+1$ as $r_k = (p_k, \hat{g}_i)$, with:

- p^k the probability of observing at least k ride requests in i during T;
- $\hat{g}_i$ the expected income from serving the corresponding potential passenger, assumed to be independent of time.

Therefore, with relocation offers, local controllers communicate to drivers the likelihood of being matched with a passenger if they relocate in the region i and the revenue drivers can expect from such a relocation.

The analysis of historical data supports the estimation of those two features. On one side, analyzing the historical requests data allows us to characterize the ride-hailing demand probability distributions and estimate p^k. Conversely, $\hat{g}_i$ can be derived from historical trip data. Based on this information, vehicles evaluate and compare their utility in relocating towards one region or the

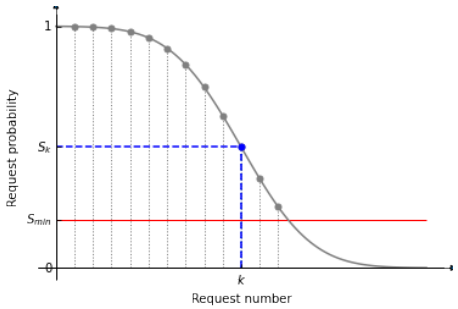

FIGURE 1: Definition of requests probability occurrences given the demand zonal distribution

other, sort their favorite relocation options and bid accordingly.

In theory, local controllers can emit an infinite number of relocation offers (each associated with a decreasing probability) and let the vehicles decide whether to apply or not on them. In practice, for realism and computational cost issues, local controllers only emit associated with an occurrence probability exceeding a minimum threshold p_{min} (set to 0.2 in this study). Figure 1 illustrates partitioning the predicted future demand into uncertain requests.

Vehicle-to-Controller negotiation scheme

The periodic auctioning process is based on the distributed Gale-Shapley algorithm[15–17]. We summarize below the outline of our algorithm:

- Local controllers predict the future local demand (*potential future passengers*), and deduce their future vehicle needs from it. They express those needs by publishing *relocating options* into a virtual two-sided *relocating market*. Each relocating option is associated with the likelihood that the demand will occur.
- With a fixed frequency, vacant vehicles enter the market and evaluate their utility for each relocating option. 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. When bidding, vehicles declare their expected arrival time at their destination.
- Local controllers agents gather the bids of vehicles and evaluate the utility expected from the vehicle service, *i.e.*, the ability of the bidding driver to relocate on time for the anticipated demand. The local controller accepts the best vehicle's bid for each anticipated passenger and rejects the others. A region having previously accepted a vehicle bid can reject it if it receives a more appealing offer at a successive matching round.

The process keeps on; rejected vehicles apply to their next favorite relocation offer until running out of interesting options. This process is illustrated in Figure 2.

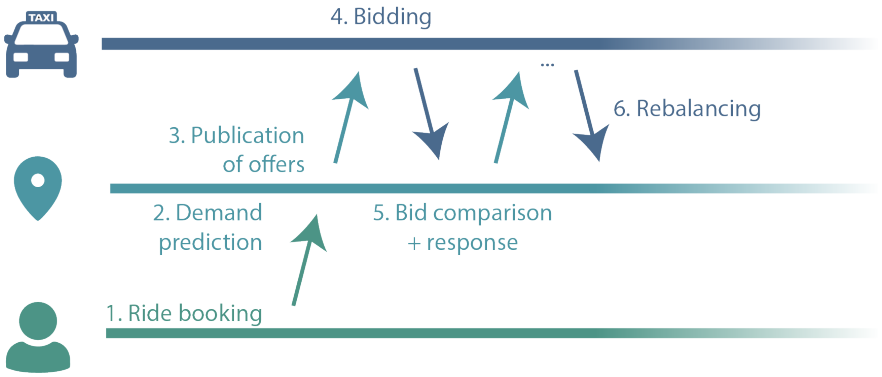

FIGURE 2: Communication protocol

Utility functions

Different utility functions can be envisioned here, but we focus at this stage on one function per agent.

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 $c_{rebalancing}(v,i)$:

$$U_v(i,k) = g(i,k) - c_{rebalancing}(v,i) \tag{1}$$

$$= \bar{g}_i \cdot p^k - c_v^{km} \cdot d_{SP}(v,i) \tag{2}$$

The expected earnings are estimated from the relocation offer features as the product of the ride-request occurrence with the average earnings when serving a passenger in the area. The rebalancing costs are computed as the product of the vehicle empty kilometric cost c_v^{km} with the relocating distance, computed as the shortest path between vehicle position v and the depot of the region i.

Utility of local controllers

Local controllers gather simultaneously and independently the applications of vehicles to the various relocation offers they published. Their utility is a decaying function of the travel time vehicles need to relocate. Therefore, given two vehicles applying to the same relocation offer, the one closest to the region depot will benefit the local controller more. On one side, it ensures the vehicle's fast arrival and quick passenger pick-up. From a more global point of view, this utility setting limits the empty mileage of vehicles.

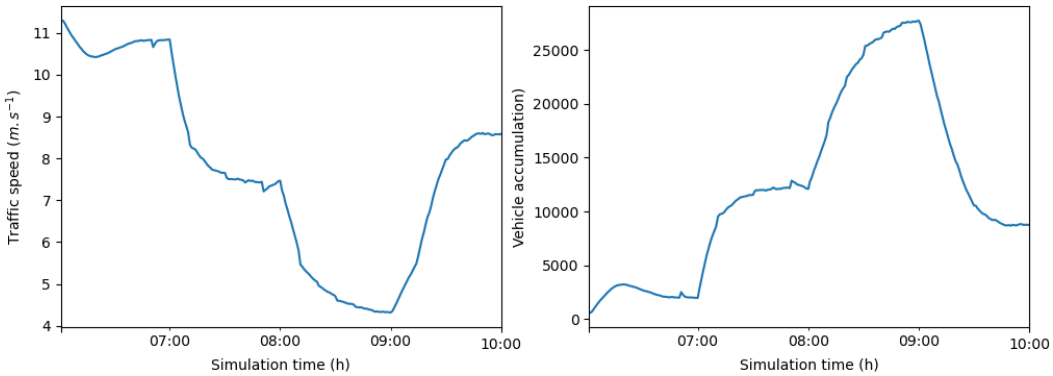

FIGURE 3: Traffic dynamics throughout a simulation

SIMULATION PLAN AND ENVIRONMENT

Simulation platform

We conduct an extensive simulation-based analysis in the open-source Python simulation platform MnMs (Multi-modal Network Modeling and Simulation) [18]. The platform relies on a multi-modal trip-based MFD (Macroscopic Fundamental Diagram) traffic model [19, 20]. It supervises the movements of vehicles and users over the network, applying the same collective mean speed per mode and urban region. Figure 3 illustrates the accumulation and dynamic speed profiles during one simulation. The platform includes a fleet management component that models different mobility services, such as individual modes, public transportation, and ride-sourcing (including ride-hailing and ride-splitting). For ride-sourcing services, the platform allows customization of default ride-sourcing management strategies addressing fleet dispatching and rebalancing issues.

Case study definition

City specific features and demand balance

The city of Lyon, France, is selected as a case study. Lyon is the third most populated city in France and the center of France's second most populated urban area. Figure 4 represents at a fine scale the neighborhoods of Lyon with inflow/outflow unbalance during the morning peak (6 am to 10 am), considering flows emitted from the city center and outer suburbs not represented on the map. This representation allows us to identify the geographic patterns of the demand during that period and anticipate the need for fleet rebalancing when operating a ride-sourcing service. 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. To a lesser extent, other regions of Lyon, such as the southern industrial neighborhoods, attract morning flows. Although these spatial demand patterns characterize private car users, they illustrate how ride-sourcing vehicles operating a similar demand are likely to get accumulated within that area if no rebalancing strategy is implemented.

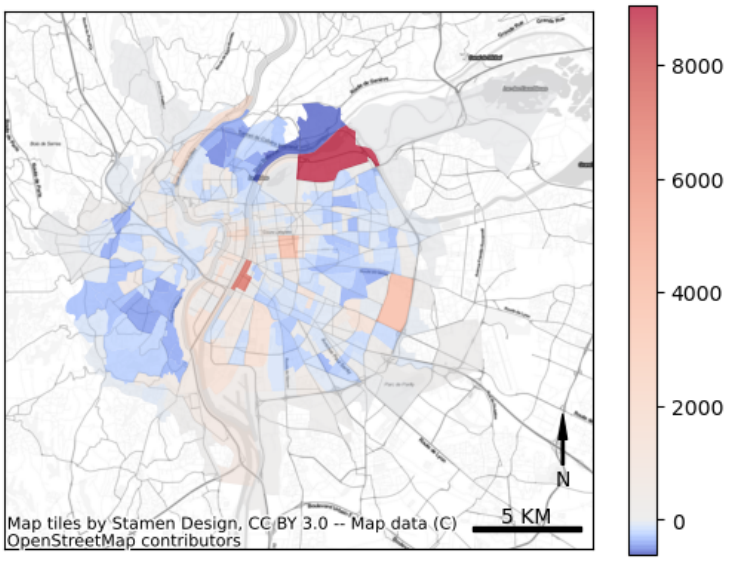

FIGURE 4: Travelers inflow and outflow during the morning peak

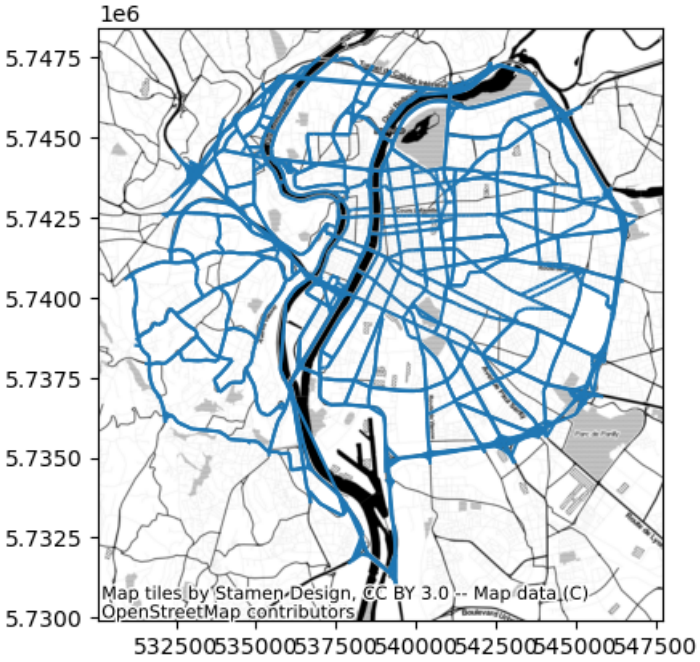

FIGURE 5: Simulation network: major road sections of Lyon, France

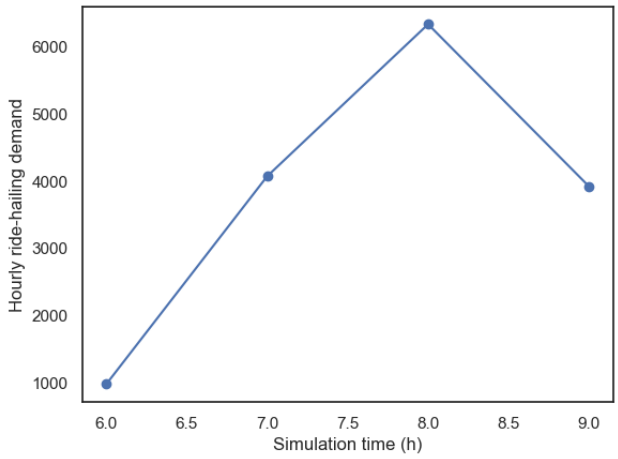

FIGURE 6: Overall ride-hailing demand

Simulated network

The simulated network covers 121 square kilometers and includes the city of Lyon and the city of Villeurbanne, both located within the main circular ring road. In order to conduct rapid simulations, we model the traffic on a simplified network of the city, made of highways, primary and secondary road sections only. The corresponding network is illustrated in Figure 5. This simplified network version comprises 5586 links and 3605 nodes for a total link length of 542 km. The supply (MFD) calibration and the demand (discussed hereafter) were previously computed within the ERC Magnum Project [21].

Demand scenario

In this work, we simulate 4 hours of road traffic, representing the 4 hours of typical morning peak, between 6 and 10 am. The demand scenario defines individuals willing to travel from an origin to a destination at a given time while prioritizing one transportation mode (private car or ride-hailing). Stochastic demand scenarios are elaborated based on a dynamic hourly origin-destination matrix describing the trips in the city of Lyon and its urban area for a typical day. The spatial granularity of the OD matrix corresponds to the IRIS (reference spatial division of the French statistical institute) zoning. Individual demand is generated by sampling origin and destination uniformly amongst nodes of the origin and destination IRIS zones and sampling departure time with a Poisson process. We additionally assign 15% of internal flows to ride-hailing mode, with the remaining users using private cars instead. The temporal hourly profile of the resulting ride-hailing demand is plotted in Figure 6.

Service areas and depots locations

We partition the city into 50 service areas to support implementing our rebalancing strategy. The partitioning is derived from a K-Means clustering of the IRIS zoning of the city. It is illustrated in Figure 7. Each service area is managed by a local controller and associated with one or several depots. We call *well-connected nodes* nodes located at road intersections with at least two different downstream directions. To study the effect of depot locations, we explore five different depot scenarios, illustrated in Figure 8, and defined as follows:

- Scenario 1: Each service area contains a single depot located at a random well-connected

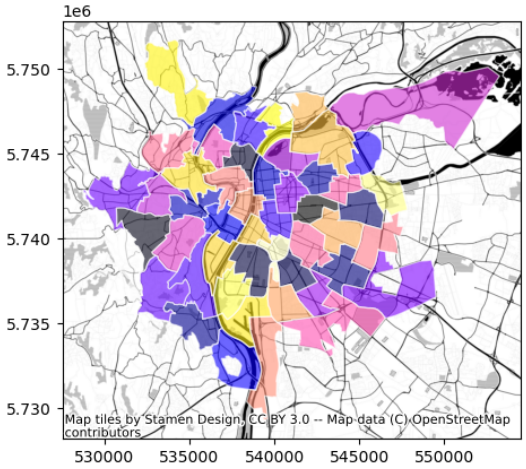

FIGURE 7: Partitioning of Lyon into 50 demand zones

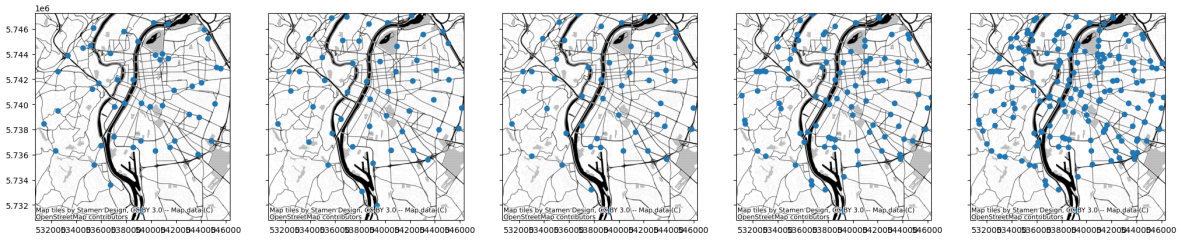

FIGURE 8: Different scenarios of depots locations

node of the area;

- Scenario 2: Each service area contains a single depot located at the centroid node of the area;
- Scenario 3: Each service area contains a single depot located at the closest well-connected node from the centroid of the area;
- Scenario 4: Each service area contains up to 3 depots, one located as in scenario 3, the others located at random well-connected nodes;
- Scenario 5: Each service area contains up to 5 depots, one located as in scenario 3, the others located at random well-connected nodes.

located at the centroid of the region.

Scenario 3 is the default depot mesh scenario used for all other sensitivity analyses.

Simulation settings

At the beginning of the simulation, all vehicles are empty and idling and uniformly distributed on the network. The auctioning and rebalancing period is fixed at 10 minutes. After rebalancing, vehicles are available for matching. Requesting passengers are matched with the nearest available vehicle.

The ride pricing scheme c_{ride} is designed as the sum of a fixed service cost $c_{service}$ and a mileage fare applied to the shortest path distance between the pick-up and drop-off point:

Parameter		Unit	Values
Ride-sourcing fleet size	M	vehicles	{1000, 2000, **3000**, 4000}
Maximum matching time	δ_{match}	minutes	1
Max. waiting time before service abandonment	$\delta_{pick-up}$	minutes	{**5**, 10}
Rebalancing frequency	Δ_r	minutes	10
Solo trip base fare	f_{base}	€	2.20
Solo trip per-km fare	f_{km}	€/km	1.00
Empty per-km cost	$c_{s,v}$	€/km	{**0.5**, 1, 2, 3, 4}
Demand prediction uncertainty	$\widetilde{\sigma}$	·	{**0.1**, 0.2, 0.3}
Request probability mean threshold	p_{min}	·	0.2

TABLE 1: Simulation parameters

$$f_{ride}(p_o, p_d) = f_{base} + f_{km} \cdot d_{SP}(p_o, p_d) \tag{3}$$

Table 1 summarizes the parameters used in this paper. Bold parameters used within parameter ranges indicate the default values.

Simulation plan and Key Performance Indicators (KPIs)

The result section successively explores the performance of the rebalancing strategy and its sensitivity to:

- different passenger patience tolerance scenarios;
- different fleet size scenarios;
- demand uncertainty and rebalancing costs;
- various depot locations strategies.

In those studies, the performance of the service is systematically analyzed through the scope of quality of service (number of passengers served and waiting time) and vehicle activities (idle time and vehicle traveled distance per served passenger).

RESULTS

Overall performances in various passengers waiting tolerance settings

Service perfomances

We compare the service performances with and without implementing the fleet rebalancing. The waiting time tolerance of passengers $\delta_{pick-up}$ is expected to impact each strategy's results significantly. Therefore, we compare the results in two settings, $\delta_{pick-up} = 5$ minutes and $\delta_{pick-up} = 10$ minutes. Table 2 gathers the performance of each strategy in terms of the number of passengers served and passengers' average waiting time before pick-up. The table also records these performances for the whole simulation time (6 am to 10 am) and the peak hour only (8 am to 9 am). We make the following observations.

- First, in both waiting tolerance settings, the rebalancing strategy successfully contributes to increasing the number of passengers served and reducing the waiting time.
- Second, the peak hour particularly benefits from the rebalancing strategy implementation, with more significant passenger increases and waiting time drops.
- Finally, the increase in passengers served is stronger for $\delta_{pick-up} = 5$ minutes (+18.41% during peak hour) than for $\delta_{pick-up} = 10$ minutes (+4.80%). Instead, the waiting time

				5 minutes				10 minutes	
		Reference	Rebalancing	Abs	Rel (%)	Reference	Rebalancing	Abs	Rel (%)
No. served passengers	Global	13,647.00	14,764.00	1,117.00	8.18	14,837.00	15,121.00	284.00	1.91
	Peak hour	4,996.00	5,916.00	920.00	18.41	5,940.00	6,225.00	285.00	4.80
Waiting time (in min.)	Global	1.50	1.01	-0.50	-33.07	2.11	1.27	-0.84	-39.90
	Peak hour	2.35	1.56	-0.79	-33.72	3.63	1.93	-1.70	-46.74

TABLE 2: Service performance comparison with different passengers waiting time tolerance

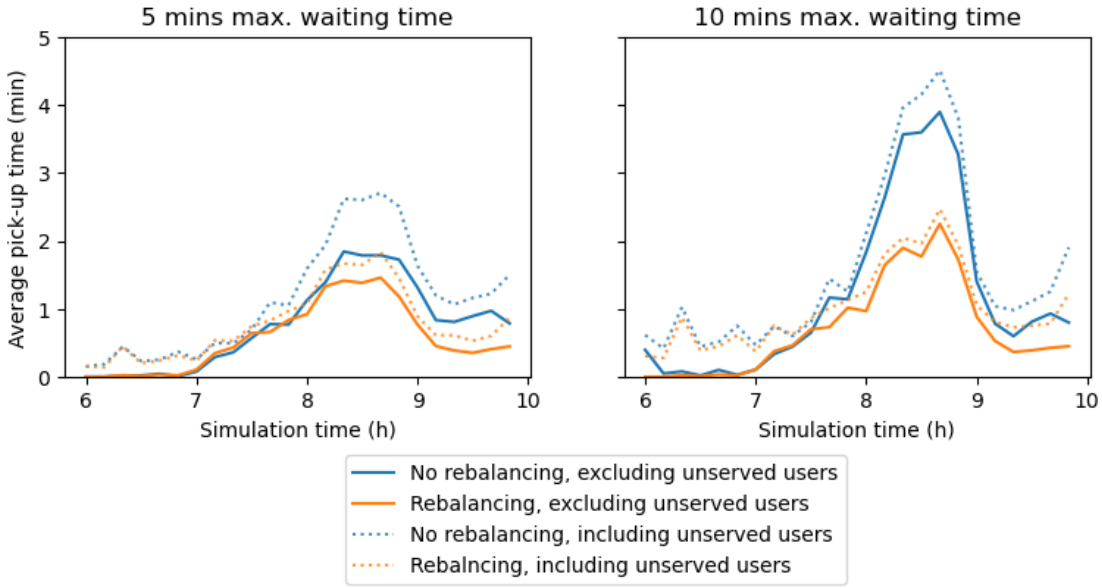

FIGURE 9: Waiting time evolution through simulation in both waiting tolerance settings

reduction is more significant if $\delta_{pick-up} = 10$ minutes (-1.70 minutes) rather than if $\delta_{pick-up} = 5$ minutes (-0.79 minutes).

This observation highlights two aspects of the benefits of a rebalancing strategy. On the one hand, under tight time constraints, fleet rebalancing enables the service to reduce abandonment by serving additional passengers. However, the pick-up time windows are such that the waiting time cannot be significantly reduced. When passengers are more patient, the service without rebalancing already meets demand satisfactorily. Nevertheless, rebalancing can still boost service quality by reducing passengers' pick-up times.

Figure 9 displays the evolution of passengers pick-up waiting throughout the simulation for each waiting time tolerance setting. The *no rebalancing* strategy is represented in blue lines, and the rebalancing strategy is in orange. The solid lines represent the average waiting time for served passengers only, while the dotted lines correspond to waiting time, including unserved passengers. The plots confirm that the passengers' waiting time reduction is especially significant when passengers are patient. In both cases, the rebalancing strategy limits the waiting time increase during the peak. Plus, considering all users, fleet rebalancing delays the waiting time increase and ensures a faster recovery.

The remaining analyses of this paper have been conducted with a waiting tolerance $\delta_{pick-up} =$ 5 minutes.

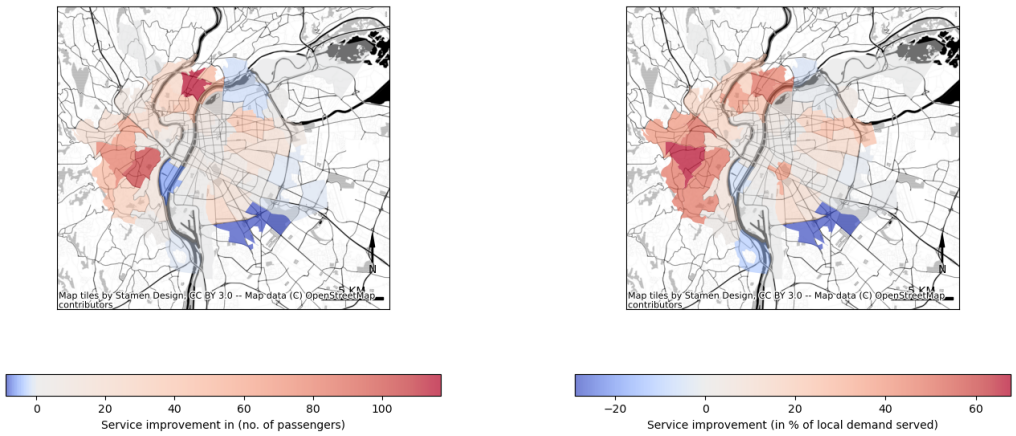

(a) Map of the service evolution in number of additional local passengers served

(b) Map of the service evolution in percentage of local demand served

FIGURE 10: Zonal service improvement in number of passengers served

Spatial performances

In this section, we analyze the service performance at a finer scale. Figure 10 illustrates the spatial distribution of the zonal increase of passengers served. Figure 10a displays its absolute evolution. Instead, Figure 10b represents the extra share of local demand served with rebalancing.

Regarding the absolute increase of demand served, the rebalancing strategy is especially beneficial to a few neighborhoods in the western and northern parts of Lyon. These neighborhoods are characterized by a reduced-density road network (cf. Figure 5) due to hilly relief and mainly residential land use. When looking at the extra share of local demand satisfied, the results are less contrasted, and many neighborhoods appear to benefit from the fleet rebalancing, allowing to serve an extra 60% of the local demand. The areas benefiting the least or suffering from the rebalancing strategy correspond to areas likely to face vehicle accumulation in Figure 4.

Passenger waiting times are also analyzed with this granularity and displayed in Figure 11. The map evidences a uniform decrease in waiting time in the western and northern areas of Lyon and the city center. Concerning waiting time, the southern service area suffers the most from the strategy implementation. As the demand in this area is minimal, the variation of service for users there strongly impacts the local service perception.

Finally, Figure 12 represents the locations of unserved passengers. The rebalancing strategy succeeds in drastically reducing the number of unserved positions. However, some service gaps subsist, with well-defined areas concentrating on the frequently unserved positions. For some neighborhoods (cf. in the southwest of the city), it is due to the distance from the nearest depots. For some others (cf. in the south-east of the city), the location of nearby depots does not prevent a high density of unserved requests. It is a consequence of working with a directed network and occurs when depots are located downstream of the demand. Then, the vehicles are relocated to the depots to serve this local demand but later fail in serving it because it implies significant detours and pick-up times. Vehicles from other nearby depots may answer the matching request instead,

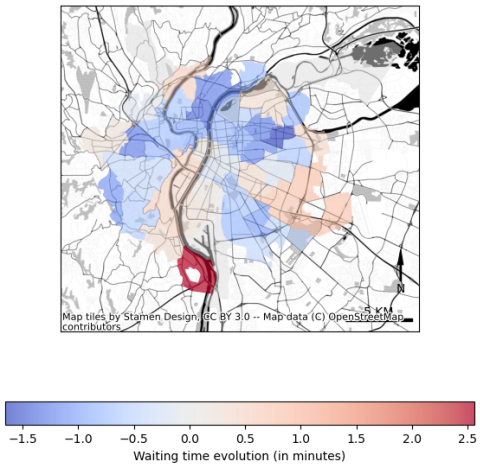

FIGURE 11: Map of zonal average waiting time reduction

leading to vehicles shortage in those depots and propagating demand dissatisfaction.

This limit of the approach and sensitivity to depot location is explored later in this paper.

Vehicles activities

Looking into the rebalancing impact on vehicle activities, Figure 13 illustrates the states of the ride-hailing fleet throughout the simulation. Compared to the *no rebalancing* strategy, the periodic rebalancing of vehicles significantly increases the number of in-service vehicles during the demand peak. However, we observe that a share of rebalancing vehicles does not systematically translate into a similar increase in serving vehicles. The demand uncertainty and situations of depots located downstream of the demand partially explain this observation, which needs further investigation. During the peak, due to the travel time increase, relocating vehicles are still traveling when the next rebalancing period start. It may result in vehicles arriving too late into the next rebalancing period and missing the passenger they were relocating for, contributing to "failed" relocations. Sensitivity to rebalancing frequency is not explored in this work, but this aspect will be studied in future work to evaluate how it alters service performances.

Table 3 summarizes the numerical values of some key indicators to evaluate the activities of vehicles, as they could be considered proxies for the drivers' revenue estimation. With no surprises, since the number of passengers served rises, the average idle time (*i.e.*, the time when vehicles are neither picking-up nor serving a passenger) of vehicles decreases significantly. However, the vehicle traveled distance per passenger increases significantly, meaning that the rebalancing strategy inflicts additional empty distance for serving passengers compared to the *no rebalancing* strategy. This result is further discussed in the next section.

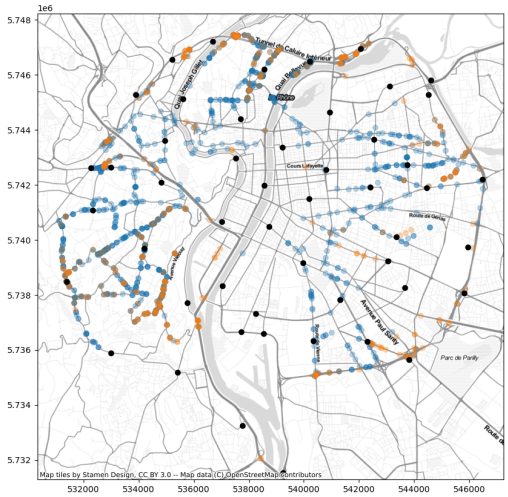

FIGURE 12: Spatial distribution of unserved users, without rebalancing (in blue) and with rebalancing (in orange). The black dots represent the depots locations.

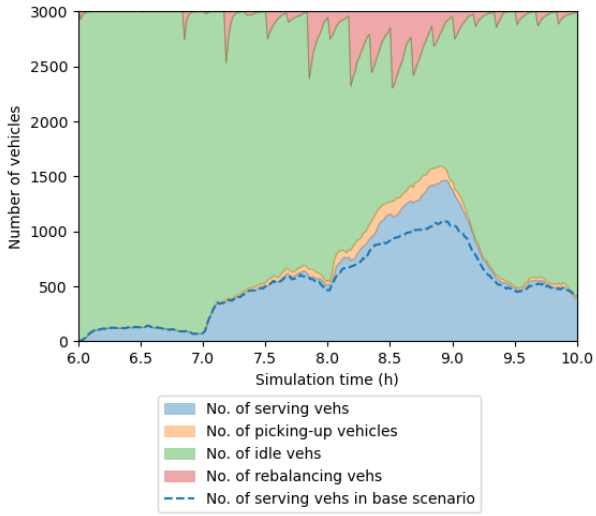

FIGURE 13: Distribution of the activities of vehicles throughout simulation time

	Reference	Rebalancing	Abs	Rel (%)
Avera idle time (min)	64.88	58.38	-6.50	-10.02
Average traveled distance per served passenger (km)	3.80	4.41	0.61	16.00

TABLE 3: Comparison of vehicles activity indicators

	Number of passengers				Waiting time (mins)			
	Reference	Rebalancing	Abs	Rel (%)	Reference	Rebalancing	Abs	Rel (%)
1000	11278.00	11891.00	613.00	5.44	2.60	2.36	-0.23	-9.02
2000	12728.00	14250.00	1522.00	11.96	1.94	1.38	-0.56	-28.93
3000	13647.00	14764.00	1117.00	8.18	1.50	1.01	-0.50	-33.07
4000	14124.00	14860.00	736.00	5.21	1.19	0.86	-0.33	-27.69

TABLE 4: Sensitivity of service performances to fleet size

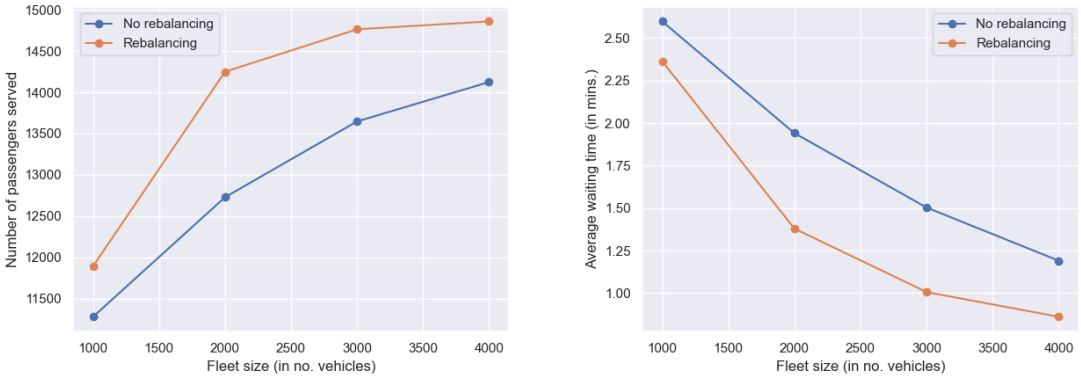

(a) Number of passengers

(b) Waiting time

FIGURE 14: Sensitivity of service quality to fleet size

Sensitivity to fleet size

Service performances

The effect of fleet size on service quality is assessed and displayed in Table 4 and Figure 14. Figure 14a illustrates the effect of the rebalancing strategy on the number of passengers served. Instead, Figure 14b displays the effect of the rebalancing strategy on the number of passengers served. The rebalancing strategy is shown to have the most significant impact with the middle fleet scenarios (2000 and 3000 vehicles). These scenarios provide a large fleet with vacant vehicles available for rebalancing but limited enough to require rebalancing. The 2000-vehicles scenario provides the highest service quality increase in passengers (+11.96%). Instead, the 3000 vehicles scenario ensures the highest waiting time reduction (-33.07%).

Vehicle activities

Extending the analysis carried on in the previous section, we look into the impacts of fleet size and rebalancing on vehicle activities. While the rebalancing strategy allows serving more passengers per vehicle, the increase in fleet size significantly counteracts this effect (Figure 15a). The activity time of vehicles follows similar trends (Figure 15b).

Looking into eniquity indicators, the rebalancing strategy is also shown to reduce inequities between vehicles. Figure 16 displays for instance the interdecile ratio $\frac{D_8}{D_2}$ and difference between $D_8 - D_1$. Except for the more extensive fleet, the rebalancing strategy reduces the relative and absolute difference between the most and least active vehicles of the fleet.

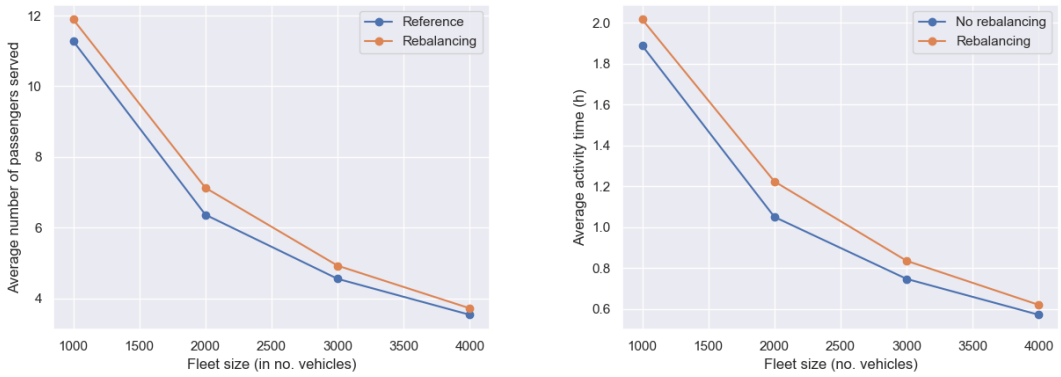

(a) Average no. passengers served per vehicle (b) Vehicle idle time reduction

FIGURE 15: Sensitivity of vehicles paid activities to fleet size

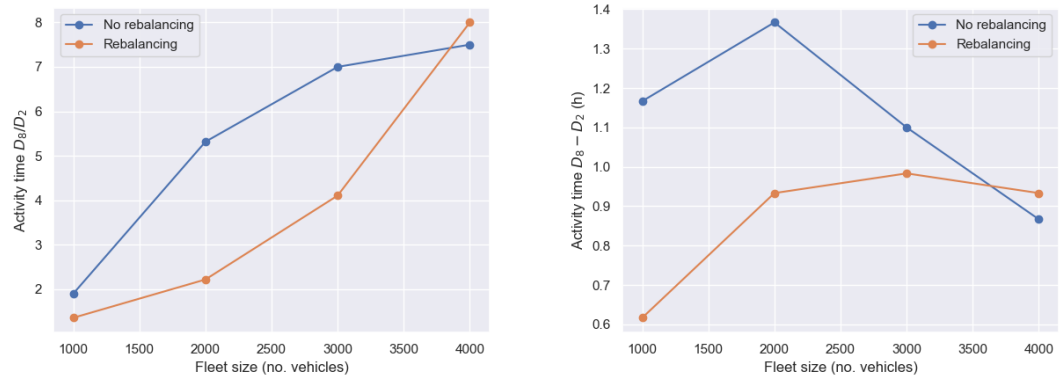

(a) Average no. passengers served per vehicle (b) Vehicle idle time reduction

FIGURE 16: Sensitivity of vehicles paid activities to fleet size

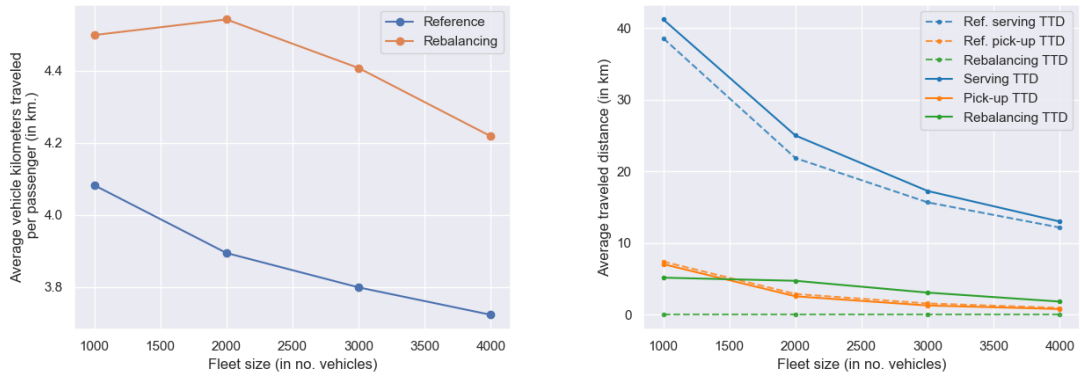

(a) Vehicle traveled distance per served passenger

(b) Contributions of service activities to average traveled distances

FIGURE 17: Evolution of traveled distances with fleet size

However, as illustrated in Figure 17a, the rebalancing strategy implies a general increase in the average distance vehicles travel per passenger. Figure 17b details the total distances traveled per activity. As the number of passengers served increased with rebalancing, the distances traveled for service also increased. The rebalancing strategy slightly decreases the pick-up distances, meaning that vehicles located at depots are, on average, closer to customers than without fleet rebalancing. However, the new rebalancing distances largely counteract this pick-up distance decrease.

Several factors can explain this significant extra travel distance:

- First, when vehicles rebalance, they head to the region depot before traveling from the depot to the passenger's pick-up location. This results in extra distances compared to a direct pick-up.
- Second, as the network is sparse and directed, the shortest paths from the vacant vehicle position to the depot and from the depot to the pick-up position imply detours and overestimated traveled distance compared to the actual network. Densifying the depots' mesh or optimizing their location can contribute to reducing this additional distance.
- Finally, possible mismatches in the rebalancing process due to demand uncertainty may result in unsuccessful rebalancing for drivers, requiring them to relocate again.

Sensitivity to demand uncertainty and rebalancing costs

This section explores the joint effect of demand uncertainty and rebalancing costs on service performances. On the one hand, the higher the demand uncertainty, the more relocation offers are published, but with reduced occurrence probabilities, making them less attractive for vehicles. On the other hand, relocating costs also impact the vehicle rebalancing utility (cf. Equation 2). We explore uncertainty levels ranging from 10 to 30% of the mean demand, and the rebalancing kilometric cost varies between 0.5€/km and 4€/km. Recall the vehicles compute the utility of relocating to region i as the difference between the expected revenue from a ride starting in i (which depends on the ride occurrence probability p^k, the base fare f_b, the mileage fare $f_k m$ and expected ride distance) and the relocation costs to i. Analytically, the ratio $\frac{f_b}{f_{km}}$ defines the maximum distance vehicles would be willing to relocate if the expected ride distance was 0 km and vehicles got

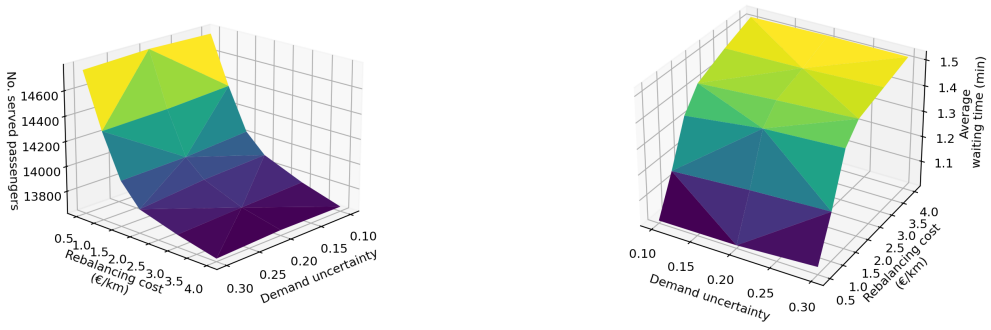

(a) Evolution of no. passengers served

(b) Evolution of passengers waiting time

FIGURE 18: Evolution of service quality with demand uncertainty and rebalancing costs

		Rebalancing cost					
		0.50	1.00	1.50	2.00	3.00	4.00
Demand uncertainty	0.1	1.01	1.16	1.33	1.42	1.47	1.52
	0.2	1.01	1.18	1.37	1.44	1.48	1.52
	0.3	1.03	1.19	1.39	1.45	1.49	1.52
	Relative increase (%)	1.94	2.82	3.98	2.48	0.95	0.11

TABLE 5: Impact of demand uncertainty and rebalancing cost on passengers average waiting time in minutes

paid with the base fare f_b only. As f_b is set to 2.2€, a rebalancing cost f_{km} varying between 0.5 and 4€/km implies that vehicles would relocate at most 0.5 to 4.4 km in such conditions.

Service performances

Figure 18 illustrates the evolution of the service quality with varying demand uncertainty and rebalancing costs. Figure 18a presents the evolution of the total number of passengers served and Figure 18b the impact on the average waiting time of passengers before pick-up. The rebalancing cost is shown to have the most significant effect on these two variables. The increased rebalancing costs significantly decrease the number of passengers served and increase their average waiting time by up to 1.5 minutes. When the rebalancing cost is particularly high, vehicles barely relocate, which explains that the performances are similar to the performance of the *no rebalancing* strategy. To a lesser extent, the demand uncertainty also decreases the performance of the service. As shown in Table 5, the demand uncertainty increases significantly when the rebalancing costs are reasonable and do not prevent vehicles from relocating. Under those circumstances, an increase of the demand uncertainty from 10 to 30% can increase the passengers waiting time by up to 4%.

Vehicle activities

However, the demand uncertainty more strongly impacts the vehicle's activities than the service quality. Figure 19 illustrates the effect of demand uncertainty and rebalancing costs on the average

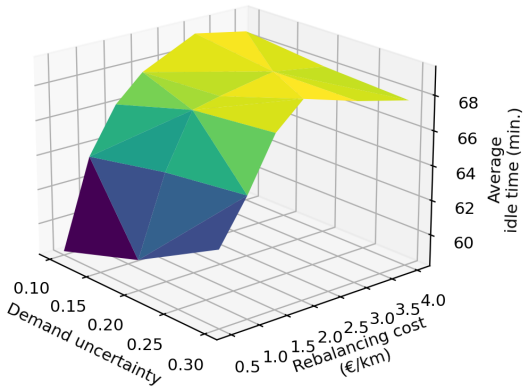

FIGURE 19: Evolution of vehicle average idle time with demand uncertainty and rebalancing costs

idle times of vehicles. Unsurprisingly, as it strongly restricts the number of relocating vehicles, the increase in rebalancing costs reduces the vehicle's activities and contributes to an increase in the average idle time of vehicles. However, when the rebalancing costs are low, we also more clearly observe the effect of demand uncertainty on average idle time. For instance, when $c_{empty} =$ 0.5€/km, an increase of the demand uncertainty from 10% to 30% contributes to a 7.12% increase in the idle time of vehicles. The vehicles are more frequently unsuccessfully rebalanced, implying more significant idle times. In future work, the effect of uncertainty will be assessed on smaller fleets, *i.e.*, in more tense operation settings. Likely, the demand uncertainty will impact not only the supply but the demand side.

Effect of deposits number and locations

In Section 5.2.2, we discussed the effect of the rebalancing strategy on the traveled distances. We suggested that the features of the network and locations of the depot could impact it. To measure this effect, this section compares the service performance of the rebalancing strategy for the five different depot location scenarios defined in Section 4.2.4.

Service performances

Figure 20 compares the performances of the rebalancing strategy for the different depot mesh scenarios. Figure 20a compares their performances regarding the number of served passengers, and Figure 20b looks into passengers' waiting time.

Focusing on scenarios 1, 2, and 3, with a single depot per service area, we observe in Figure 20a that locating them on well-connected nodes (as in scenarios 1 and 3) rather than on the centroid node (scenario 2) increases the number of passengers served. Additionally, picking the well-connected node close to the centroid (as in scenario 3) rather than a random well-connected node (scenario 1) also improves the number of served passengers. Selecting a central depot (scenarios 2 and 3) also contributes to reducing the passengers' waiting time (cf. Figure 20b).

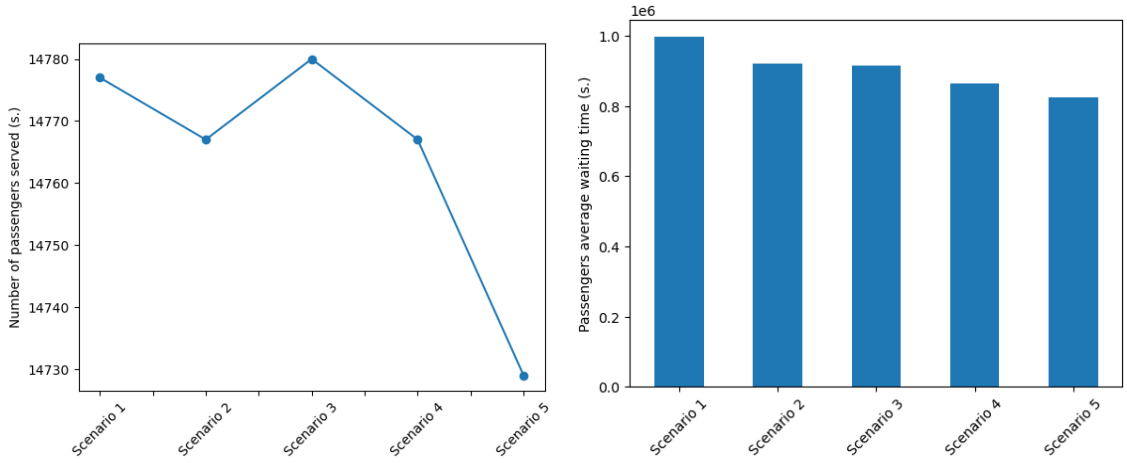

(a) Number of passengers served for each depot mesh scenario

(b) Passengers average waiting time for each depot mesh scenario

FIGURE 20: Service performances with different depot location strategies

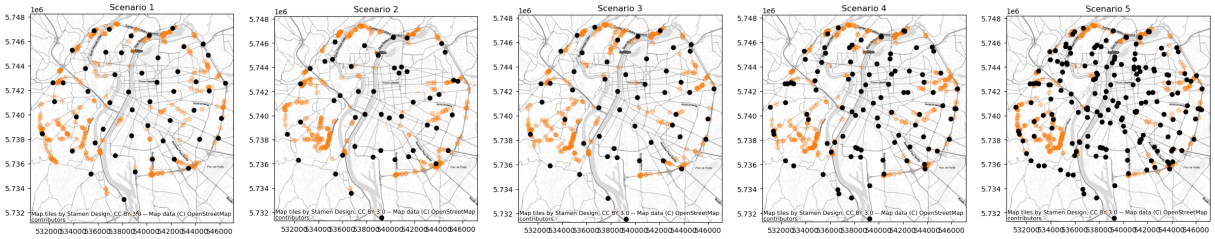

FIGURE 21: Positions of unserved users due to exceeded pick-up time.

Interestingly, randomly adding extra depots to Scenario 3, as in Scenarios 4 and 5, does not allow to increase the number of passengers served. These results should be further investigated, mainly by generating larger sets of dense networks, but it suggests that a single central depot position better serves the zonal demand. However, this densification reduces the average waiting time of passengers because it still drastically reduces the waiting time of many passengers.

For each scenario, Figure 21 illustrates where passengers exceeding their pickup waiting time are located on the network. We observe that most depot locations satisfy the nearby local demand but that the depot densification is not enough if depots are located downstream of the flow. Therefore, the depot positioning strategy should better account for the depot topology and road section directions.

Locating depots on well-connected nodes implies that the demand along corridors or isolated roads suffers from lower satisfaction. In Scenarios 2 to 5, a neighborhood in the southwest part of Lyon is especially poorly served. The depots of these zones are located quite far away, on an outer ring road, and distances to reach this neighborhood from there is too large to satisfy the local demand.

Finally, Figure 22 confirms that the location of depots substantially impacts the vehicle traveled distances per served passenger. Although it does not permit satisfying the same number of passengers, locating the depot at the region centroid is the scenario with one depot, which limits

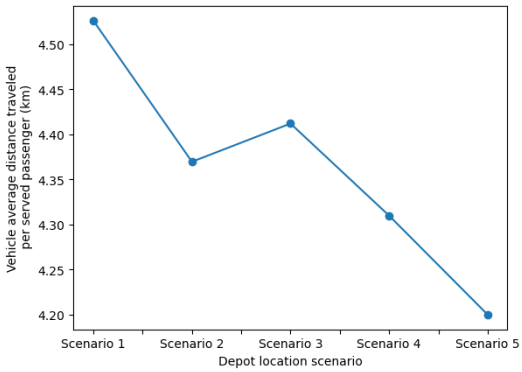

FIGURE 22: Caption

the most traveled distance per served passenger. Moreover, as anticipated, increasing the number of depots per region (Scenario 4 and 5) does contribute to reducing the vehicle kilometers traveled per passenger.

This final study confirms that the service performances and externalities are sensitive to the depot mesh design. Central and well-connected nodes offer the best performances but do not suffice to satisfy the demand in some corridors and areas with little accessibility. Densifying the depot network reduces passengers waiting time and extra kilometers traveled, but locating these additional depots more strategically seems necessary to ensure a better network and demand coverage.

CONCLUSION

Through extensive simulation studies, this paper investigates the performance of an original auctioning process for large-scale ride-hailing fleet rebalancing. The process involves outsourcing relocating directions from the service to local controllers who negotiate with idle vehicles based on hypothetical future demand. The distributed Gale-Shapley algorithm is employed to match idle vehicles with relocation offers.

The strategy is applied in the city of Lyon, France, using the trip-based MFD simulation platform MnMS for fast computing and individual tracking of vehicles and travelers. Simulations cover four hours of complete morning peak demand, with 15% of the inner demand prioritizing ride-hailing over private cars. To understand the strategy's performances, the simulations explore various scenarios, including passenger waiting tolerance, fleet size, demand uncertainty, rebalancing costs, and depot locations.

Results demonstrate a significant impact on the number of passengers served and waiting time, particularly during demand peaks. Spatial analyses reveal benefits in the suburbs with limited accessibility and network density. Moreover, the rebalancing strategy ensures a more equitable distribution of paid activity time among drivers, as indicated by two different inequity metrics.

However, the method shows increased distance traveled per passenger served, meaning additional contribution to road traffic and its externalities. The paper addresses this issue by evaluating different depot location strategies. While a more robust mesh design is needed, this analysis provides guidelines for improving and developing it.

As a continuation of this work, future research will focus on better assessing the effect of uncertainty, exploring non-uniform uncertainty distribution, and developing local incentive strate-

gies to encourage vehicle relocation to areas with lower accessibility or uncertain demand. Additionally, the performance of the service in ride-splitting and multi-provider operating modes will be assessed.

AKNOWLEDGMENTS

This project has received funding from the European Union's Horizon 2020 research and innovation program under Grant Agreement no. 953783 (DIT4TraM).

REFERENCES

1. Beojone, C. and N. Geroliminis, On the inefficiency of ride-sourcing services towards urban congestion. *Transportation Research Part C: Emerging Technologies*, Vol. 124, 2021, p. 102890.
2. Hryhoryeva, M. and L. Leclercq, Competition in Ride-Hailing Service Operations: Impacts on Travel Distances and Service Performance. *Transporation Research Record*, 2023.
3. Wallar, A., 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)*, 2018, pp. 4539–4546, iSSN: 2153-0866.
4. Simonetto, A., J. Monteil, and C. Gambella, Real-time city-scale ridesharing via linear assignment problems. *Transportation Research Part C: Emerging Technologies*, Vol. 101, 2019, pp. 208–232.
5. Liu, Y. and S. Samaranayake, Proactive Rebalancing and Speed-Up Techniques for On-Demand High Capacity Ridesourcing Services. *IEEE Transactions on Intelligent Transportation Systems*, 2020, pp. 1–8.
6. Alonso-Mora, J., 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)*, 2017, pp. 3583–3590.
7. Miao, F., S. Han, S. Lin, Q. Wang, J. Stankovic, A. Hendawi, D. Zhang, T. He, and G. J. Pappas, Data-Driven Robust Taxi Dispatch under Demand Uncertainties. *arXiv:1603.06263 [cs]*, 2017.
8. Miao, F., S. Han, S. Lin, J. A. Stankovic, H. Huang, D. Zhang, S. Munir, T. He, and G. J. Pappas, 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*, 2015, pp. 100–109.
9. Ramezani, M. and M. Nourinejad, Dynamic modeling and control of taxi services in large-scale urban networks: A macroscopic approach. *Transportation Research Part C*, Vol. 94, 2018, pp. 203–219.
10. 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.
11. Zhu, P., I. I. Sirmatel, G. Trecate, and N. Geroliminis, Distributed Coverage Control for Vehicle Rebalancing in Mobility-on-Demand Systems. *TRB Annual Meeting*, 2022, pp. 22–03340.
12. M. Seppecher and L. Leclercq, A Decentralised Auction-Based Framework for Rebalanc-

ing Ride-Hailing Fleets with Uncertain Request Probabilities. In *102nd Transportation Research Board Annual Meeting (TRB 2023)*, Washington D.C., USA, 2023.

13. Manjunath, A., 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*, 2021, pp. 1–17.
14. Wu, Y. H., 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* (S. Nittel, A. Labrinidis, and A. Stefanidis, eds.), Springer, Berlin, Heidelberg, Lecture Notes in Computer Science, 2008, pp. 252–270.
15. Gale, D. and L. S. Shapley, College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, Vol. 69, No. 1, 1962, pp. 9–15.
16. Brito, I. and P. Meseguer, Distributed Stable Matching Problems with Ties and Incomplete Lists. In *Principles and Practice of Constraint Programming - CP 2006* (D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, and F. Benhamou, eds.), Springer Berlin Heidelberg, Berlin, Heidelberg, Vol. 4204, 2006, pp. 675–679.
17. Ostrovsky, R. and W. Rosenbaum, Fast Distributed Almost Stable Matchings. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing*, ACM, Donostia-San Sebastián Spain, 2015, pp. 101–108.
18. LICIT-ECO7 Lab, Univ. Gustave Eiffel, Lyon, France, *MnMS (Multimodal Network Modelling and Simulation)*, 2023.
19. Mariotte, G., L. Leclercq, and J. A. Laval, Macroscopic Urban Dynamics: Analytical and Numerical Comparisons of Existing Models. *Transportation Research Part B: Methodological*, Vol. 101, 2017, pp. 245–267.
20. Leclercq, L., A. Sénécat, and G. Mariotte, Dynamic Macroscopic Simulation of On-Street Parking Search: A Trip-Based Approach. *Transportation Research Part B: Methodological*, Vol. 101, 2017, pp. 268–282.
21. Mariotte, G., 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, 2020, pp. 62–86.