

Effective Operation of Demand-Responsive Transport (DRT): Implementation and Evaluation of Various Rebalancing Strategies

Chengqi Lu^{1*}, Dr. Michal Maciejewski¹, Prof. Dr. Kai Nagel¹

Institut für Land- und Seeverkehr, TU Berlin, Germany

*Corresponding Author: lu@vsp.tu-berlin.de

Abstract

In this study, four rebalancing strategies are implemented in an open-source multi-agent traffic simulation platform, MATSim. The impact of the rebalancing strategies on service quality of the Demand-Responsive Transport (DRT) is examined and compared. A comprehensive evaluation scheme, consisting of both conventional and innovative criteria, is introduced for this purpose. Results indicate that the claimed effectiveness of all the rebalancing strategies included in this study can be materialized in the agent-based simulation platform to certain degrees. Depending on the role a DRT system plays, different strategies may be the preferred choice. In addition, as the by-product of this study, the structure update of MATSim and the introduction of new evaluation criteria enable convenient implementation and evaluation of custom rebalancing strategies.

Keywords:

Demand-Responsive Transport and Fleet Rebalancing

Introduction

Demand-Responsive Transport (DRT) is a mobility-on-demand service, where customers can submit their travel demand spontaneously and the operator will try to send a vehicle to fulfil that demand. Nowadays, DRT services at different scales can be found in many places around the world. Most of them are operated by Transportation Network Companies or conventional taxi companies. Recently, there have also been several studies that propose the idea of viewing DRT as (part of) the public transport service provided by the city government or transport authority (Wang, 2018), (Sieber, Ruch, Horl, Axhausen, & Frazzoli, 2019) and (Vakayil, Gruel, & Samaranayake, 2017). When we consider DRT as a public service, instead of a commercial one provided by profit-making companies, the way to operate the DRT system can be different from those we are often seeing now. Rather than putting profitability at the top place, service quality, accessibility as well as the impact on environment and traffic congestion will probably move up the rank in the priority list of the designer and operator of the system.

The operation of the DRT service mainly consists of the following two parts: assignment strategy and rebalancing strategy. As suggested by its name, the assignment strategy matches the available vehicles and travel requests. The rebalancing strategy repositions the DRT fleet such that there will be enough vehicles in every region and the travel requests can be served shortly after the submission. In this study, we focus on the rebalancing strategy in the DRT operation. We implement several existing rebalancing

strategies into MATSim, a multi-agent transport simulation platform. Then we will evaluate and compare the impact of the rebalancing strategies on the service quality.

About the MATSim simulation platform

Multi-Agent Transport Simulation (MATSim) is an open-source agent-based framework for traffic simulations (W. Axhausen, Horni, & Nagel, 2016). Because of its efficient design, city-scale traffic simulation can be conducted. There are multiple extension modules in the MATSim library, focusing on modelling various components or aspects of transport systems. In this study, we add new rebalancing strategies to the DRT extension module (Bischoff, Maciejewski, & Nagel, 2017). The MATSim DRT module enables the simulation of the on-demand transport service with the help of within day planning feature in MATSim. At the departure time, the customer will walk to the closest link (i.e., the technical term for road segment in MATSim) to the departure location (e.g., home, work). The DRT system will then try to arrange an available vehicle to pick up the passenger. After arriving at the link closest to the destination point, the vehicle will drop off the passenger and the travel request is considered to be completed. Within this framework, fleet operational strategies can be implemented to provide a good service quality at a possibly low cost.

As the by-product of this study, the newly implemented rebalancing strategies as well as the evaluation criteria have been added to MATSim's DRT extension and made available via a configuration file. Implementation and evaluation of additional custom rebalancing strategies can also be performed conveniently. More details on this can be found in the supplementary materials.

Implementation of Rebalancing Strategies

In this section, we will briefly introduce the rebalancing strategies used in this study. As this study focuses on the implementation and the comparison of the rebalancing strategies, we will not go into the details of each strategy. Interested readers are advised to refer to the original literature.

Background information

Before introducing the rebalancing strategies used in this study, we will first clarify two concepts which can be found in some of the rebalancing strategies: zonal aggregation model and reliance on previous data (i.e., data-driven). Because of these features, it was necessary to add some new elements to the MATSim simulation environment, in order to make the rebalancing strategies work properly.

Zonal aggregation model divides the whole service area into smaller zones. This feature enables the use of optimization tools when generating rebalancing plans, because the zonal aggregation greatly reduces the problem size by condensing multiple links into one zone. Zonal aggregation model also allows computing indicators used in decision making and which may be hard to obtain at the level of single links or coordinates. When a zone-based (i.e., based on zonal aggregation model) rebalancing strategy is used, rebalancing actions are first planned and calculated at the zone level, instead of the link or

Effective Operation of Demand-Responsive-Transport (DRT)

coordinate level. Then, we convert the zonal level rebalancing plans into link-to-link rebalancing instructions that are sent to individual vehicles. There are many ways to generate zonal system and to convert zonal rebalancing plans into individual vehicle rebalancing instructions. In this study, we generate the zonal system by applying k-mean clustering technique on requests' departure locations. More details on the zonal system setup can be found in the *Scenario and simulation setup* section below.

A rebalancing that relies on previous data strategy generates rebalancing plans based on the historical demand data. Such strategies are also often referred to as data-driven strategies in the literature. As pointed out by many studies (Zhang, Rossi, & Pavone, 2016) and (Song, Kanasugi, & Shibasaki, 2016), by analysing the previous data on how people move around in the traffic network, predictions on the future trips can be made in an accurate manner. With the prediction of potential DRT trips, we can proactively relocate vehicles closer to zones for which we expect a higher demand in the near future.

Min-Cost-Flow Rebalancing Strategy

The Min-Cost-Flow Rebalancing Strategy, introduced in (Bischoff & Maciejewski, 2020) is an intuitive and straightforward rebalancing strategy that stocks vehicles in the popular zones based on the past travel data. A day is also divided into multiple time bins. At the beginning of each time bin, selected available vehicles will be sent across the network such that the target value of available vehicles for each zone could be met when there are enough available vehicles. The target value of available vehicles in a zone during a time bin is based on the historical demand data. The mapping from the historical data to the target value requires manual parametrisation. Zones with more idling vehicles than the target value will send vehicles to the zones where the target is not met. Rebalancing is modelled as the transportation problem and the interzonal rebalancing plan is calculated using the Hungarian algorithm (Ford Jr & Fulkerson, 1956), such that the total cost of the rebalancing drives is minimized.

Adaptive Real-time Rebalancing Strategy

The Adaptive Real-time Rebalancing Strategy, originally proposed in paper from (Pavone, Smith, Frazzoli, & Rus, 2012), distributes available (idling) vehicles in the DRT system evenly across the network periodically, such that vehicles will be available at any place in the system. At each rebalancing period, the strategy will count the number of available vehicles in the system. The target number of vehicles a zone should possess is calculated by dividing the total number of available vehicles by the number of zones. Similar to the Min-Cost-Flow Rebalancing Strategy, this strategy also uses the zonal aggregation model and the interzonal rebalancing plan is calculated by solving the transportation problem. But unlike the Min-Cost-Flow strategy, this strategy does not rely on the previous travel data.

Feedforward Rebalancing Strategy

The Feedforward Rebalancing Strategy, also proposed in the paper from (Pavone, Smith, Frazzoli, & Rus, 2012), models the travel demand in the DRT system as fluidic flow. The rebalancing plan is then generated by calculating the optimal counter flow to the demand flow such that the system stays balanced. It is a data-

Effective Operation of Demand-Responsive-Transport (DRT)

driven and zone-based rebalancing strategy. The service area is divided into zones and a day is divided into small time bins. During each time bin, we assume the travel demands are entering the system at a constant rate. With this framework, we can compute the optimal counter flow of the travel demands based on past data. And the optimal counter flow will be used as the feedforward signal when generating the rebalancing plan.

To better adapt this strategy to the iterative approach in MATSim, a feedback mechanism is added to this strategy. With the feedback mechanism, the rebalancing strategy will maintain a minimum number of vehicles in each zone. The feedback mechanism has a higher priority than the feedforward part. If a zone does not have spare vehicles (i.e., in addition to the minimum requirement), no vehicle will be sent out from that zone even if there is a feedforward signal instructing this zone to send out vehicles.

Plus-One Rebalancing Strategy

The Plus-One Rebalancing Strategy is a novel and model-free rebalancing strategy (Ruch, Gachter, Hakenberg, & Frazzoli, 2020). It uses the real-time travel demand as the only input to determine the way to relocate available vehicles in the system. As suggested by its name, this rebalancing strategy will relocate one idle vehicle to the departure location of each request that has been matched during the last rebalancing period. In other words, this strategy relies on the logic that a request is likely to appear in a currently popular departure area and therefore it is a good idea to send rebalancing vehicles to that area. This is particularly the case during the peak hours where commuters travel from residential areas to work or the other way round. A more formal proof on the sturdiness of this Rebalancing strategy can be found in (Ruch, Gachter, Hakenberg, & Frazzoli, 2020).

Summary

The four rebalancing strategies included in this study is summarized in Table 1. The features introduced above are included in the table to differentiate the strategies. If a strategy is zone-based, then the network needs to be pre-processed to generate the zonal system. If a strategy is data-driven, then historical travel demand data is needed.

Table 1 - Summary of the Rebalancing Strategies

	Abbreviation	Zone-based	Data-driven
Min-Cost-Flow Rebalancing Strategy	Min-Cost-Flow	x	x
Adaptive Real Time Rebalancing Strategy	Adaptive	x	
Feedforward Rebalancing Strategy	Feedforward	x	x
Plus-One Rebalancing Strategy	Plus-One		

Simulation Setup and Evaluation Criteria

Scenario and simulation setup

In this study, we use a MATSim model of Vulkaneifel to simulate and compare the rebalancing strategies. Vulkaneifel is a sparsely populated area consisting of many scattered small cities, towns and villages in

Effective Operation of Demand-Responsive-Transport (DRT)

Rhineland-Palatinate, Germany. Because of the geographical pattern, the public transport service in that area is very limited. As it is prohibitively expensive to operate a high frequency public transport service in this region, a DRT service is an attractive alternative to improve the mobility of the residents in the area.

To generate the DRT demand, we synthesize the population in Vulkaneifel based on the data provided by Senozon Deutschland GmbH with the method introduced in (Neumann & Balmer, 2020). Then the MATSim iterative approach is performed to determine the potential users of the DRT service if this service is introduced to the area (Kaddoura & Schlenther, 2020). The corresponding trips will serve as the base DRT demand profile in this study.

Because some of rebalancing strategies are data-driven, we need to include a day-to-day fluctuation to simulate the imperfectness of the demand prediction (i.e., the difference between predicted travel demand based on the historical data and the actual demand). We generate the actual DRT demand by slightly perturbing the base DRT demand profile. Specifically, we remove 10% of the DRT trips and replace them with a similar number of trips converted from other modes (e.g., car, public transport, walk and bike) within the DRT service area. Furthermore, we randomly mutate the departure time of 10% of the DRT demand by up to ± 30 minutes. For the reactive strategies, we will use the actual DRT demand directly for the simulation. For data-driven ones, we feed in the base DRT demand to the strategies and actual DRT demand will be used for simulation.

After the perturbation, there are 9694 DRT requests that need to be served by the DRT fleet (i.e. no rejection is allowed). The fleet size is variable in this study and range from 300 to 700. The fleet size to demand ratio is within the range of other studies on the DRT service in the literature (such as in (Bischoff, Maciejewski, & Nagel, 2017), (Hörl, Ruch, Becker, Frazzoli, & Axhausen, 2019), (Vosooghi, Puchinger, Jankovic, & Vouillon, 2019) and (Fagnant & Kockelman, 2016)). To match the available DRT vehicles to the travel request, a simple matching strategy is used. The matching strategy is first come first serve and will assign the closest available vehicle to the request. Pooling is not enabled, so each vehicle can serve only one request at a time.

The zonal system in this study is generated with k-mean clustering technique. First, the k-mean cluster algorithm is run on the base DRT demand profile. We run the clustering algorithm based on requests' departure locations. After several trial runs, setting the number of clusters to 110 will yield a good zone size for the Vulkaneifel scenario. Then we use the central coordinates (i.e. the mean coordinates of all the points within the cluster) of each cluster to generate Voronoi polygons. By intersecting the Voronoi polygons with the service area, the DRT zonal system is generated. Finally, some manual adjustment is performed for the zones near the boundary of the service area, such that there are no disconnected zones or very small zones. Figure 1 shows the DRT zonal system on the map. The brown dots represent departure locations of DRT requests in the base demand profile. The red dots are central

coordinates of the requests within each cluster.

In order to convert a zone-level rebalance plan into vehicle level instructions, we apply a simple rule. For a rebalance action $R(X, Y, m)$, which means m vehicles should be sent from zone X to zone Y , we will randomly choose m vehicles from zone X , and send them to the road segment in zone Y that is closest to the central coordinates of the requests within zone Y (i.e. the red dot in Figure 1).



Figure 1 – Zonal system based on request clustering

Evaluation criteria

In order to compare the performance of different rebalancing strategies, we need to come up with a set of evaluation criteria. In this study, we not only include the commonly used criteria based on waiting time and travel distance statistics, but also introduce new criteria to evaluate the rebalancing strategies. We evaluate the service quality of the DRT by comparing the following indicators: mean waiting time, median waiting time, 95 percentile waiting time, empty driving distance, and zonal vehicle availability. The first four criteria indicate the overall service quality and efficiency of the system. The 95 percentile waiting time also reflects the guarantee of the minimum service level. The zonal vehicle availability, which is a newly developed criterion in this study, focuses on the service level experienced by the users at different places. Combining all these evaluation criteria, a comprehensive evaluation of the DRT system can be achieved.

The reason to use 95 percentile waiting time value instead of the maximum waiting time to show the minimum service level is because the former statistic is much more stable. The maximum value is oversensitive to the outlier effect and that can make the comparison among different rebalancing strategies and different simulation setups less persuasive. The 95 percentile, on the other hand, is much less sensible to the outliers. And still, it provides a good indication for the minimum service level of the system. As DRT is considered a service that responds to spontaneous travel demand, the 95 percentile waiting time of the DRT system should be comparable to a high frequency public transport service.

Effective Operation of Demand-Responsive-Transport (DRT)

The zonal vehicle availability rate refers to the percentage of time during which there is at least one available vehicle in that zone or in one of the neighbouring zones. The availability $A(z)$ of the zone z is calculated through equations below, where T is the total simulation time (in second, which is the basic simulation time unit in MATSim), $V(z, t)$ is the number of available vehicles inside zone z at time t and $N(z)$ is the set of the neighbouring zones of z .

$$A(z) = \frac{\sum_{t=0}^T \max(\alpha(z, t), 0.5 * \beta(z, t))}{T}$$

$$\alpha(z, t) = \begin{cases} 1, & V(z, t) \geq 1 \\ 0, & \text{Otherwise} \end{cases}$$

$$\beta(z, t) = \begin{cases} 1, & \sum_{y \in N(z)} V(y, t) \geq 1 \\ 0, & \text{Otherwise} \end{cases}$$

With the zonal availability criterion, we can determine whether the service is readily available at a region in the network. This criterion works as a compliment to the zonal waiting time statistics. For zones with low travel demand, the zonal waiting time statistics does not reveal the service quality for that region because of the lack of samples. The availability, on the other hand, can still serve as a good indication for the service quality for those zones.

Results and Analysis

The obtained mean, median and 95 percentile waiting times and empty travel distances are shown in Figure 2 (a) – (d). Since pooling is not allowed, the travel distance with customer onboard is the same for all the cases. For a clearer comparison, we therefore use the empty travel distance here. For some rebalancing strategies, the line does not cover the full fleet size range of 300 to 700. This is because those strategies do not work properly for a small fleet and the DRT system is no longer able to serve all the requests. As a result, those setups are considered as infeasible and are not included in the plots.

From the results, we can observe that all the rebalancing strategies reduce the overall waiting time at the expense of extra travel distance. Among the rebalancing strategies, the two proactive ones (Min-Cost-Flow and Feedforward), the reduction in the mean waiting time is more significant compared to the other two strategies. The difference in the median waiting times among the strategies are slightly smaller. Given the adequate fleet size, the Feedforward rebalancing strategy provides lower 95 percentile waiting time. When we investigate the total travel distance, the Min-Cost-Flow strategy is the clear winner, with only a very slight increase in distance compared to the case without rebalancing. The use of Adaptive rebalancing strategy, on the other hand, results in the largest increase in total travel distance.

Effective Operation of Demand-Responsive-Transport (DRT)

The reason behind the Feedforward and Plus-One rebalancing strategies not working properly for smaller fleets is their algorithmic design. The rebalancing will not be turned off even if there are not enough vehicles to serve the travel demands. It may happen that all the available vehicles left are being relocated and there will be no idle vehicle to serve the incoming requests, and consequently the DRT system gradually loses the stability (i.e., the number of unserved requests keeps building up until the end of the service hour), which results in a very bad performance in the end. This can be solved by adding simple empirical techniques. As this study focuses on the performance of rebalancing strategies under normal working conditions, a temporal disabling of rebalancing strategies is not investigated.

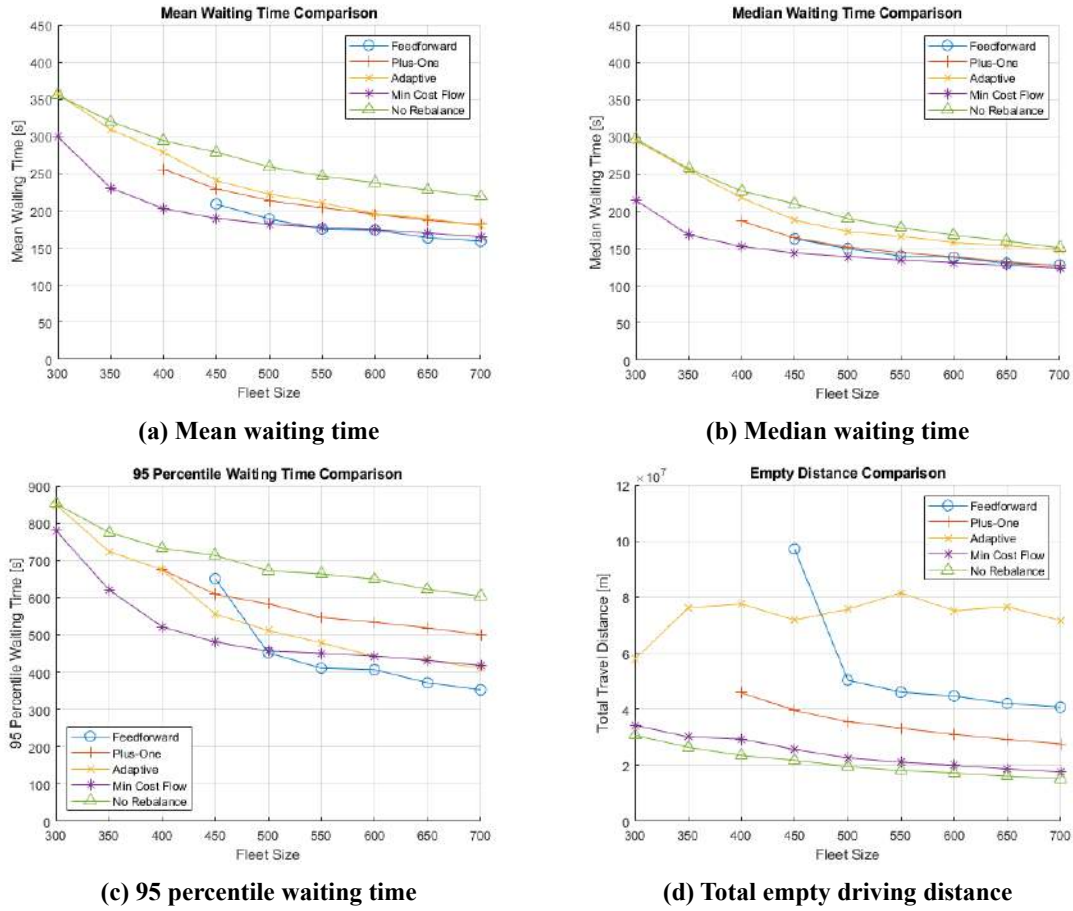


Figure 2 – System-wide results

Based on Figure 2, we can claim that all the rebalancing strategies work properly with 500 vehicles and any further increase in the fleet size gives only a slight improvement in service quality. Therefore, we present the whole-day spatial vehicle availability only for the fleet size of 500 in Figure 3. As mentioned above, the zonal availability rate serves as a complimentary evaluation criteria to the waiting time statistics. To provide readers a full image of the service quality experienced by the customers in different locations, the zonal level 95 percentile waiting time is also included (see Figure 4). It needs to be pointed out that because of the uneven distribution of the travel demands, some regions may have very few requests and 95 percentile waiting time may be susceptible to the slight changes in the travel demand.

Effective Operation of Demand-Responsive-Transport (DRT)

Therefore, we include the illustration of the level of demand in different zones in Figure 5. For zones with very few demands, the 95 percentile statistics should be referred to with caution. The results obtained for other fleet sizes can be found in the supplementary material.

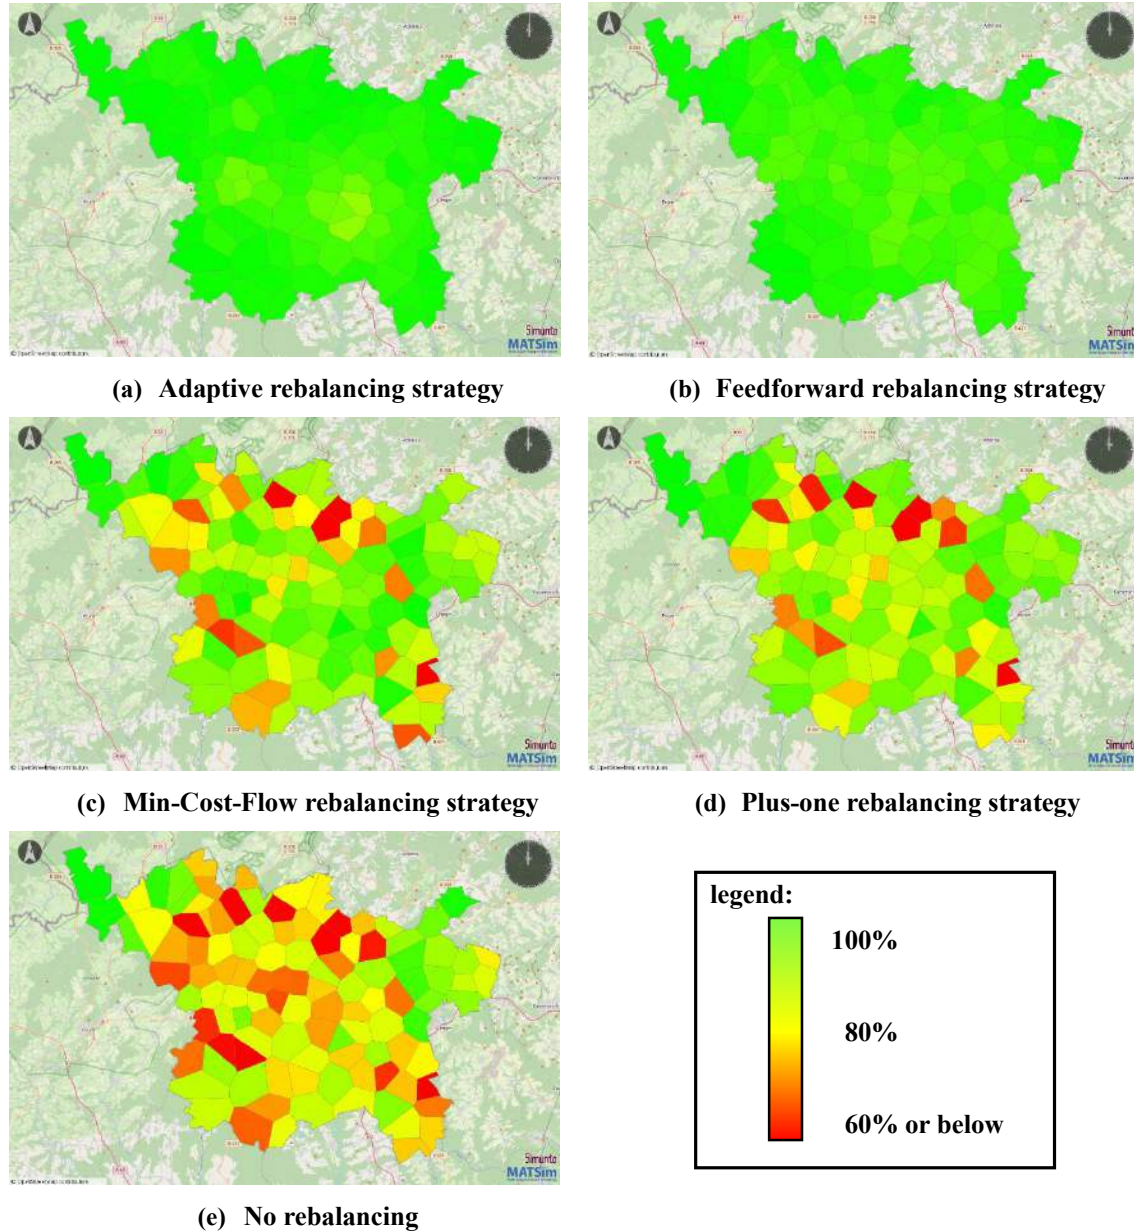


Figure 3 – Availability illustration for fleet size of 500

Comparing plots presented in Figure 3 and Figure 4, we can observe that Feedforward and Adaptive rebalancing strategies achieve a very good availability with almost all zones having the availability over 90%. That means at almost any place in the DRT service area, there will be at least one vehicle readily available nearby, for over 90% of the time throughout the day. Meanwhile, the 95 percentile waiting time in most of the zones are also below 10 minutes. For the Min-Cost-Flow and Plus One rebalancing strategies, the low availability in some zones is because the number of vehicles relocated to a zone by

Effective Operation of Demand-Responsive-Transport (DRT)

the strategy is proportional to the number of requests originating in that area. This leads to long waiting times in low-demand zones. This effect is also illustrated by the 95 percentile waiting time in those zones (e.g., the zones in the bottom left in Figure 4 (c) and (d)). Nevertheless, compared to the no rebalance case, the overall availability and 95 percentile waiting time throughout the network is still improved to some extent.

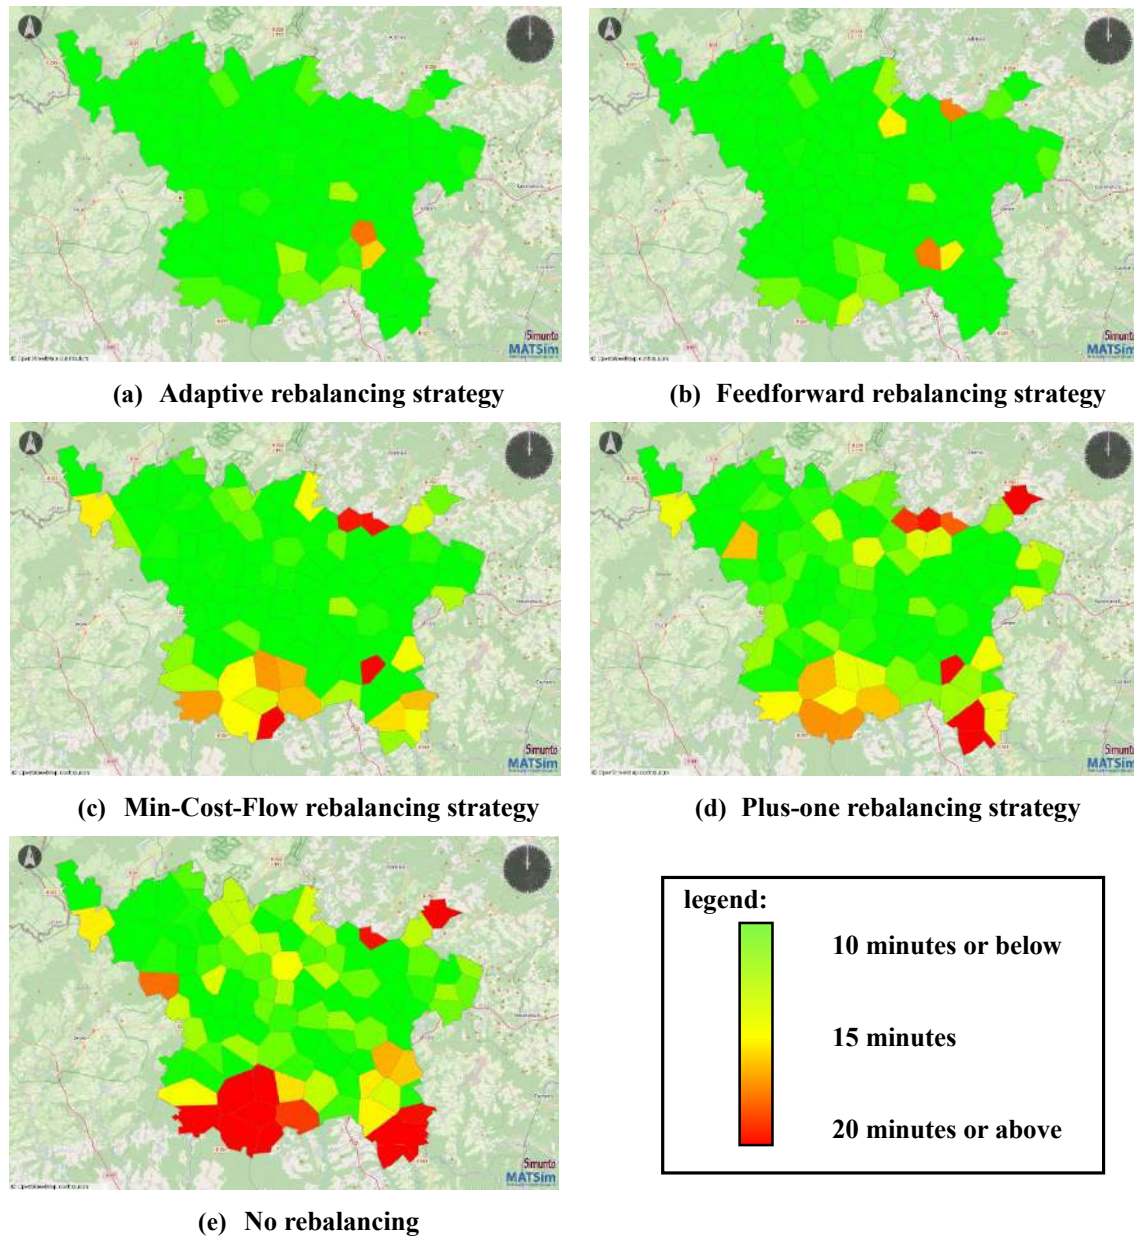


Figure 4 - Zonal 95-percentile waiting time illustration for fleet size of 500

If we use the plots in Figure 2 as the only evaluation criteria, which is a common approach in the existing literature, then the Min-Cost-Flow rebalancing strategy will be the best strategy. If a DRT operator has to run the service with a small fleet, then the Min-Cost-Flow will be the only choice that can improve the service quality. When the fleet size is adequately large, both Feedforward and Min-Cost-Flow

Effective Operation of Demand-Responsive-Transport (DRT)

rebalancing strategies can provide a significant improvement in service quality, but the Min-Cost-Flow strategy achieves that at a smaller cost in terms of the total empty drive distance. Therefore, applying the Min-Cost-Flow rebalancing strategy will be the preferred strategy. This result is to some extent surprising, as the Min-Cost-Flow strategy is a relatively simple and intuitive rebalancing strategy. Despite its simplicity, it provides similar or even better performance than its more sophisticated counterparts.




Figure 5 - Zonal demand level illustration

When we also include availability analysis and zonal level 95 percentile waiting time into our scope, the results may be a little bit different. Equipped with a feedback mechanism that tries to maintain a minimum number of vehicles in each zone, the feedforward strategy provides the best availability across the whole network. The adaptive rebalancing strategy, which spreads available vehicles evenly throughout the network also achieves a very good result in terms of availability. The rest rebalancing strategies only provide good availability in certain areas. If a customer submits a request from a less popular region, a long waiting time may be experienced, especially during the peak hours, where many vehicles are relocated to the area with more customers.

Conclusion and Outlook

Based on the simulation results, it can be confirmed that all the rebalancing strategies covered in this study can improve the service quality to certain extent. Extra travel distance, on the other hand, is to be expected. With the help of the rebalancing strategies, almost all the requests can be served within 10 minutes in Vulkaneifel scenario, even with a relatively small fleet size. This makes the DRT a truly spontaneous service, where no timetable or pre-booking is necessary.

In addition, we also observe that the two data-driven strategies implemented in this study achieve better performance, even if there is day-to-day fluctuation in the travel demands. In particular, relying on the previous travel data, the intuitive Min-Cost-Flow rebalancing strategy provides a very good performance

Effective Operation of Demand-Responsive-Transport (DRT)
in terms of waiting time statistics and total travel distance.

When availability analysis is included, which can be an equally important criteria when we view DRT as a mobility service to the public, the Feedforward strategy and adaptive rebalancing strategy show superiority over the other strategies. This is no magic though, as the algorithm design of these two rebalancing strategies take the guarantee of service quality all over the service area into consideration. Yet, there is no free lunch. If we want to have high availability throughout the network, we not only need to have a relatively large fleet size, extra travel distance should also be expected.

Finally, based on the outcome and by-product of this study, we have identified several further research topics. With the update on the programming structure of the DRT extension in MATSim, rebalancing strategy can be easily simulated in different scenarios. We will carry out more simulations in various scenarios. For example, the rebalancing strategies can be examined in the city setup, where the travel demands are much heavier and densely distributed. In addition, the congestion problem is another interesting topic for further studies. Another interesting research topic is combining rebalancing with pooled rides, where the relation between demand and supply is more complex. Last but not the least, there is currently no mechanism in any of the rebalancing strategy that takes human behaviour (i.e., mode choice evolution, reaction to the fare) into consideration and the results may converge to an unrealistic equilibrium when the ordinary MATSim iterative approach is used. Therefore, along with the matching algorithm, taking the users' reaction into account is another direction for further studies.

References

1. Bischoff, J., & Maciejewski, M. (2020). Proactive empty vehicle rebalancing for Demand Responsive Transport services. *Procedia Computer Science*, 739-744.
2. Bischoff, J., Maciejewski, M., & Nagel, K. (2017). City-wide shared taxis: A simulation study in Berlin. *IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*.
3. Fagnant, D. J., & Kockelman, K. M. (2016). Dynamic ride-sharing and fleet sizing for a system of shared autonomous vehicles in Austin, Texas. *Transportation*, 45, 143–158.
4. Ford Jr, L. R., & Fulkerson, D. R. (1956). Solving the transportation problem. *Management Science*, 3(1), 24-32.
5. Furth, P. G., & Wilson, N. H. (1981). Setting Frequencies on Bus Routes: Theory and Practice. *60th Annual Meeting of the Transportation Research Board*.
6. Hörl, S., Ruch, C., Becker, F., Frazzoli, E., & Axhausen, K. W. (2019). Fleet operational policies for automated mobility: A simulation assessment for Zurich. *Transportation Research Part C: Emerging Technologies*, 102, 20-31.
7. Kaddoura, I., & Schlenther, T. (2020). *The impact of trip density on the fleet size and pooling rate of ride-hailing services: A simulation study*. Working Paper, TU Berlin.
8. Neumann, A., & Balmer, M. (2020). *Mobility pattern recognition (mpr) und anonymisierung von mobilfunkdaten*. White paper, Senozon Deutschland GmbH and Senozon AG.
9. Pavone, M., Smith, S. L., Frazzoli, E., & Rus, D. (2012). Robotic load balancing for mobility-on-demand systems. *The International Journal of Robotics Research*, 31(7), 839-854.
10. Ruch, C., Gachter, J., Hakenberg, J., & Frazzoli, E. (2020). The+ 1 Method: Model-Free Adaptive Repositioning Policies for Robotic Multi-Agent Systems. *IEEE Transactions on Network Science and Engineering*, 7(4), 3171 - 3184.
11. Sieber, L., Ruch, C., Horl, S., Axhausen, K. W., & Frazzoli, E. (2019). *Improved public transportation in rural areas with self-driving cars: The Example of Swiss Train Lines*. Research Collection, ETH Zurich.
12. Song, X., Kanasugi, H., & Shibasaki, R. (2016). Deeptransport: Prediction and simulation of human mobility and transportation mode at a citywide level. *International Joint Conference on Artificial Intelligence*.
13. Vakayil, A., Gruel, W., & Samaranayake, S. (2017). Integrating Shared-Vehicle Mobility-on-Demand Systems with Public Transit. *Transportation Research Board 96th Annual Meeting*. Washington DC, United States.
14. Vosooghi, R., Puchinger, J., Jankovic, M., & Vouillon, A. (2019). Shared Autonomous Vehicle Simulation and Service Design. *Transportation Research Part C: Emerging Technologies*, 15-33.
15. W. Axhausen, K., Horni, A., & Nagel, K. (2016). *The Multi-Agent Transport Simulation MATSim*. London: Ubiquity Press. DOI: <http://dx.doi.org/10.5334/baw>. License: CC-BY 4.0.
16. Wang, B. (2018). Simulation of Autonomous Transit On Demand for Fleet Size and Deployment Strategy Optimization. *Procedia Computer Science*, Volume 130, Pages 797-802.

17. Zhang, R., Rossi, F., & Pavone, M. (2016). Model predictive control of autonomous mobility-on-demand systems. *2016 IEEE International Conference on Robotics and Automation (ICRA)*.