Microsimulation of an Autonomous Taxi-System in Munich

Florian Dandl, Benedikt Bracher, Klaus Bogenberger
Munich University of the Federal Armed Forces, Institute for Intelligent Transportation Systems,
85577 Neubiberg, Germany, E-Mail: florian.dandl@unibw.de

***Abstract*—Studies about autonomous taxis (aTaxis) concluded that empty vehicle movements account for a share of about 10% of vehicle miles traveled. In many of these studies, constant (in time or space) travel times are used due to a lack of data and to simplify the computations. Furthermore, the influence of the empty vehicle movements on the street network has not yet been a focal point of research.**

To address these two issues, the operation of an aTaxi system is implemented into an existing traffic microsimulation model. On the one hand, a microsimulation depicts travel times more accurately than constant link-level values. On the other hand, the taxi movements influence the flows along the network links and thereby have the potential to change travel times in the street network.

Based on a calibrated Aimsun model of the city of Munich, a small number of scenarios are simulated: Starting from the calibrated OD matrix, a share of 10% of trips originating and ending inside the highway belt of Munich are completely or partially served by aTaxis.

In case of a 1-to-1 substitution of private trips with aTaxi requests, the network-wide delay of private vehicles only increases by 1% due to induced empty rides.

Furthermore, the differences between a simulation using link-level travel times and a traffic microsimulation are studied. Delays due to left turns and traffic lights are present in the microsimulation. Results show, that fleet operation algorithms need to address these issues, which occur in reality.

I. Introduction

The introduction of fully autonomous vehicles is likely to accelerate the growth of mobility on demand services. Both car-sharing, as well as taxi (or transportation networking) companies will gain major benefits from the new technology. Automated empty rides can be used to rebalance the fleet in autonomous car-sharing while the main cost factor, namely the driver, can be saved in taxi operations.

The new mobility services have many positive aspects. The biggest potential can be achieved if their clients are people who previously used their private vehicles. Private vehicles are parked for most of the day and one autonomous taxi or car-sharing vehicle can replace the trips of many private vehicles. Additionally, people, who abolish their private vehicles, will not use these mobility services for every trip if public transportation is cheaper. This statement is supported by a study by Elliot and Shaheen [1], who depicted that introducing a car-sharing vehicle actually reduces the total amount of mileage.

However, empty rides are induced for rebalancing vehicles between customers. While the amount of extra traveled kilometers is an indicator, it is not the decisive quantity to asses the impact of rebalancing. The set of routes, which have to handle the extra load, determine the consequences of this extra empty mileage. Therefore, city planners are interested in the change of travel or delay times in the whole (street) network.

This work utilizes an existing traffic microsimulation to examine these quantities. Moreover, this approach allows the study of fleet operation algorithms under more realistic traffic situations than link-level travel times. In a nutshell, this work has the following two research objectives:

- examine the impact of empty travel on the street network
- investigate the influences of more realistic traffic conditions on the fleet performance

The remainder of this paper is organized as follows: we present the state of the art regarding traffic congestion of autonomous taxi and car-sharing systems in the next section. In Section III we present the used fleet operation model and its interface with a traffic microsimulation in Aimsun. After that, the street network and the simulated scenarios are introduced. The results of these simulations are depicted in Section V. In Section VI, we compare the fleet performance to two models with link-level travel times. Finally, we draw conclusions and discuss extensions to the model.

II. State of the Art

Shared autonomous vehicles and autonomous taxis have been studied extensively in recent years. Burns et al. [2] used an agent-based model, which was simulated with daily-averaged travel times, to show that affordable and convenient mobility can be achieved without private vehicle ownership.

Fagnant et al. [3] implemented a virtual shared autonomous fleet in Austin, Texas. Hourly changing link-level travel times from a base simulation without the shared vehicles were used to reflect different traffic conditions during a day. They concluded that a vehicle replacement rate of approximately 9.34 can be achieved. The empty trips amounted to 8% of the total fleet mileage. Finally, they showed that in spite of these induced trips, the environmental impact will probably be positive because of new, environment-friendly vehicle models with a high vehicle turnover rate and less cold starts.

Many studies focused on fleet operations like customer-vehicle assignment, rebalancing and fleet size (e.g. [4], [5] and [6]). In these works, travel times were independent of the induced traffic. A few scientific analyses also considered the feedback effect of higher travel times due to the extra utilization of the street network by autonomous taxis.

978-1-5090-6484-7/17/$31.00 ©2017 IEEE

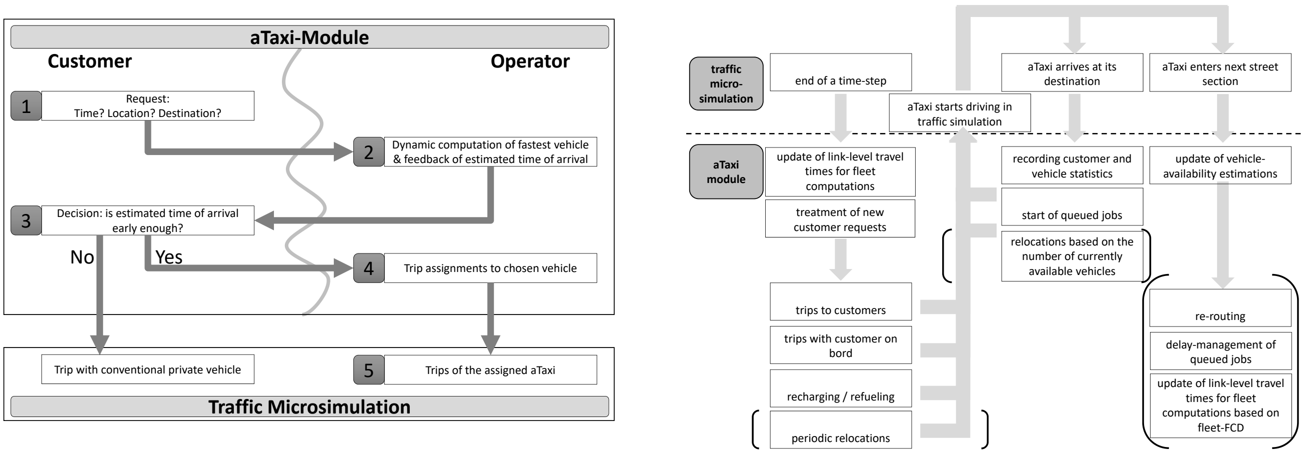

(a) Actions triggered by an incoming customer request.

(b) Interface of the aTaxi module and the traffic microsimulation. Actions in brackets are not included in this work.

Fig. 1. Schematic of fleet operation.

A framework considering this effect in a cell transmission model was presented by Levin et al. [7]. The authors focused on the case where only automated vehicles are allowed to estimate the influence of a reservation-based intersection model. Results showed that travel times could be halved.

Rossi et al. [8] conducted research about congestion-aware routing for rebalancing trips. They illustrated that in symmetric networks and the absence of privately owned vehicles it is theoretically possible to route rebalancing vehicles without generating any additional congestion. For their proof, they modeled customer trips as continuous flows in the network. This is a mathematical simplification of real customer routing, which presents itself as an integer problem. As an integer version of their method is not directly implementable, a congestion-aware algorithm was introduced, which was tested in a modified version of New York. Capacities were adjusted to be of the dimension of customer movements, and travel times depended on the aTaxi flows. It was shown that the rebalancing trips caused an increase in travel times of 2.12%.

Recently, Maciejewski and Bischoff [9] submitted a paper investigating congestion in scenarios, where both privately owned vehicles and an autonomous taxi fleet fill the street network of Berlin at the same time. Using an extended MatSim model, three scenarios about the development of street capacities in the presence of autonomous vehicles were tested, namely the capacity of nowadays, and capacity improvements of 50% and 100%. Furthermore, 6 different trip replacement rates (from 0 to 100% in steps of 20%) were assumed. All scenarios with capacity improvements resulted in lower delay times in the network. However, the scenario without capacity improvements showed rather large increases in delay times and delay ratios. The network of Berlin is not symmetric, but the drastic increase of the delay ratio from 6.5% without autonomous taxis to 15.7% without private vehicles is not in unison with the work by Rossi et al. [8].

III. Autonomous Taxi Model in a Traffic Microsimulation

This work addresses the congestion issue in yet another way. Using a calibrated microscopic traffic model as basis, a real-time executable fleet-operation algorithm is implemented in an agent-based simulation. Thereby, we test both the impact of the induced empty trips on the congestion in the network and the implications of the congestion on the fleet performance. Furthermore, the microscopic approach includes traffic lights and delays due to turns. Therefore, it represents a more realistic model - on the cost of extra computation time.

As touched on in Section II, there are quite a few papers about the operation of an autonomous fleet and multiple possibilities to control the fleet. Since this work does not focus on the operation strategy, we decided to use a rather simple model.

The traffic microsimulation is executed in Aimsun, which provides a Python API and controls the timestep. A standard value of 0.8 seconds is used for this work. A request dataset with the information $r_i = (t_i, x_i^o, x_i^d)$ is prepared in advance, where t_i is the time of the request, and x_i^o and x_i^d are the origin and destination of the potential customer, respectively. A request r_i becomes visible to the fleet operator, when the simulation time is larger than t_i. Fig. 1a schematically displays the actions that are triggered by an incoming request. The requests are treated dynamically, i.e. one-by-one in the order they appear to the fleet operator. The alternative would be a global assignment after aggregating requests for some time period. The dynamic approach has the advantage of a direct response to a customer while the global optimization will generate less customer waiting time in total. In this model, the customers decide according to the estimated waiting time computed by the fleet operator, if they use the aTaxi service.

The fleet is rebalanced only based on actual requests ("pick-up trips"), not because of demand estimations ("relocation

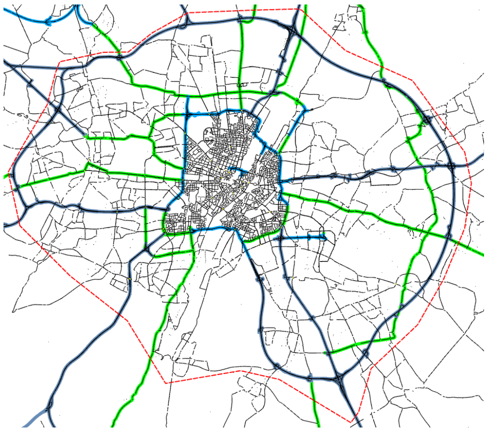

Fig. 2. Street network of the greater Munich area.

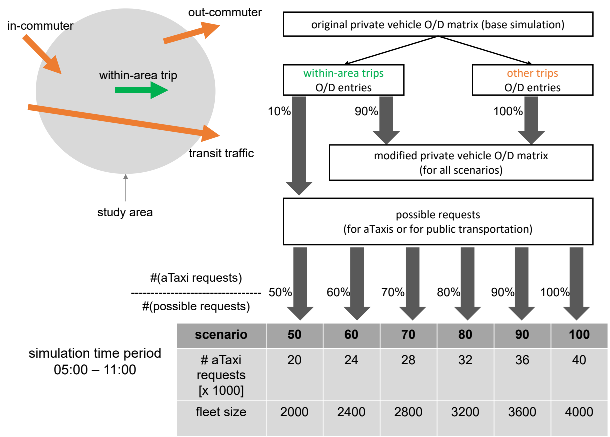

Fig. 3. Definition of scenarios.

trips"). The addition of a good relocation strategy improves customer waiting times, usually on the cost of some extra empty mileage.

The operation algorithm tracks all fleet vehicles to estimate future vehicle availability. This approach allows a vehicle-customer assignment, which is obviously better than the approach in [9], where only currently available vehicles are considered as long as there are vehicles without a job.

The model incorporates boarding times of one minute per boarding process. Boarding and deboarding can be performed simultaneously.

Fig. 1b illustrates the interactions of Aimsun and the aTaxi module. Moreover, a few possible model extensions are highlighted as well and marked by brackets.

It should be noted that even though the routes are computed by the fleet operation algorithm, OD-vehicles are created in the second row of the figure, i.e. only the information of origin and destination is transferred to the traffic simulator. Therefore, the actual routing is performed in Aimsun in this work. Since Aimsun uses the same statistics for routing as the fleet operation, the effect is small. Nevertheless, this inconsistency will be changed in a future extension.

IV. Simulations in the Greater Munich Area

A. Microscopic Traffic Model

The Aimsun microsimulation model was set up to include the region limited by the federal highway A99 and spans an area of approximately 30 x 25 km^2. The A99 is an incomplete ring highway, surrounding the city of Munich. Inside this area, the arterial network is provided with correct lane numbers, capacities, and speed limits. In the city center, which is encircled by the inner ring road B2R, the network was refined and contains every road in the area. An overview of the simulation network is given in Fig. 2.

A special focus was put on the correct intersection geometry and traffic signal configuration during the creation of the model. For 70 intersections, the real signal plans could be obtained from the city of Munich, and were included as a fixed-time schedule. For 510 other intersections, the signal plans were estimated suiting the traffic demand.

The data of 612 loop detectors, located at 174 intersections in the main road network, were used to calibrate the simulation. Because the detectors are mostly located right in front of stop lines, even single turnings can be identified. The data provided are aggregated into 15 minute intervals. With this interval being quite large compared to the cycle time of the traffic lights, the influence of traffic signals on the count data can be considered negligible. To ensure the data are representative for an average working day, only data from Tuesdays to Thursdays were used to calibrate the OD-matrix.

In Munich, the morning peak is shorter but steeper than the late-afternoon peak. Because a microsimulation is rather time-intensive, we decided to limit the simulation time on the period between 5 am and 11 am.

B. Definition of Scenarios

The uncertainty of the model, which probably has one of the largest impacts on the results, is the demand for an aTaxi system. We assume that an aTaxi system will have a limited operating area similar to free-floating carsharing nowadays. Customers will probably be able to call the service outside of the operating area as well, but for a larger fee. We therefore assume, that this will be the exception and statistically not significant for this study.

We focus on scenarios, where a share of private vehicle trips is replaced. Because of the assumption of a limited operating area, we separate the trips, which both originate and end in the study area, from the rest of the OD-matrix. The amount of these within-area trips is reduced to 90%, while the rest of the matrix entries (representing commuters and transit traffic) is left unchanged. The 10% of trips removed from the OD matrix build the set of possible requests.

Moreover, we assume that the by-trip payments of aTaxis cause people to perceive the actual cost of using a car, and thereby change their mobility behavior. As public transportation will still be cheaper than aTaxis, a share of the trips

TABLE I

TRAFFIC STATISTICS: MEAN VALUE AND STANDARD DEVIATION FROM THREE REPLICATIONS PER SCENARIO

Scenario	Private Vehicle Mileage $[10^3 km]$	aTaxi Mileage $[10^3 km]$	Ratio of Empty Mileage $[\%]$	Average Fleet Velocity $[km/h]$	Average Private Vehicle Delay $[s/km]$
Base	7958 (8)	-	-	-	74.3 (0.4)
100	7715 (19)	303 (1)	9.6 (0.1)	24.2 (0.1)	75.1 (0.9)
90	7700 (10)	271 (1)	9.8 (0.1)	24.3 (0.1)	74.1 (1.6)
80	7682 (19)	245 (2)	10.1 (0.0)	24.6 (0.1)	72.4 (0.2)
70	7669 (4)	213 (0)	10.4 (0.1)	24.9 (0.1)	72.0 (0.3)
60	7678 (14)	185 (0)	10.7 (0.1)	25.0 (0.1)	70.3 (0.5)
50	7677 (19)	155 (1)	11.2 (0.1)	25.0 (0.1)	71.9 (1.6)

previously made by car are assumed to be conducted by public transportation.

Fig. 3 summarizes the creation of 6 scenarios. The share of trips using aTaxis from the set of possible requests varies between 50% and 100%. The rest of the possible requests are considered being served by public transportation, or are covered by bike or walking.

The aTaxi fleet size is as important as the demand for the fleet performance. To allow comparability between scenarios, we scale the fleet size according to the number of requests that should be served. The rate of 10 trips per fleet vehicle is chosen out of convenience and follows the results from [3]. In reality, fleet size will be the result of an optimization process, where fleet performance pushes and financial cost per vehicle pulls the vehicle number.

V. RESULTS FROM SIMULATIONS IN AIMSUN

The indicators to measure the impact of the extra traveled empty kilometers are not that clear. Because of the different OD matrices in the base and the aTaxi scenarios and the stochastic nature of trip creation in the microsimulation, the total traveled kilometers (sums of 2nd column and the complete aTaxi mileage in table I) are not constant for the scenarios.

The percentage of empty traveled kilometers from the total fleet mileage is decreasing with an increase of the scenario scale. The area a single vehicle has to cover is larger for scenarios with smaller fleets and therefore the average approach to the customer is longer.

We decided to use the average aTaxi velocity and the average delay time per kilometer of all private vehicles in the network to quantify the traffic impacts. The velocities of fleet vehicles increase by 3% when half of the requests decide to use public transportation, cycle or walk. To evaluate the consequences of the empty mileage in case of a 1-to-1 replacement of private vehicle to aTaxi trips, we compare the average delay time of private vehicles in the complete network in the base scenario (unmodified OD matrix, no aTaxis), which is 74.3 s/km, to the value 75.1 s/km for scenario 100. The relative difference of 1% is very small. Evaluating the delay times of all vehicles (including the aTaxis and trucks), the values actually decrease, even from the base (74.1 s/km) to scenario 100 (73.9 s/km). Hence, the quantity shown in

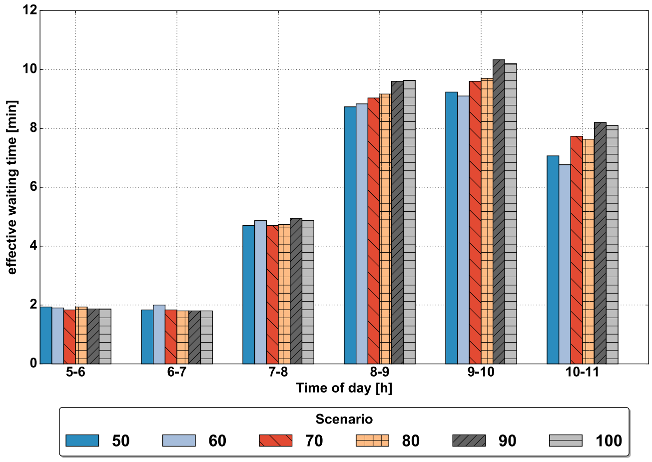

Fig. 4. Hourly averaged effective customer waiting times.

[9], total delay time over total travel time ratio, is actually decreasing in all scenarios (62-63%) compared to the base simulation (64%). These unexpected results can be explained by the aTaxis having to drive more time on side roads to approach customers. These streets have lower speed limits than the main streets and are therefore less congested. Even though the speed is not faster than on the main street, the difference to the speed limit and hence the delay time are smaller.

Next, we want to evaluate the fleet performance in the different scenarios. Because of customers having the option to deny the offered service, it is necessary to evaluate the customer waiting times and the ratio of unserved customers. Both quantities can be combined to one performance measure by the definition of an effective waiting time for request r_i:

$$T_i^{eff} = \begin{cases} T_i^{wait} & \forall r_i \text{ being served} \\ T_p & \forall r_i \text{ being unserved} \end{cases}$$

where T_i^{wait} is the actual waiting time of a served customer and T_p is a large penalty time rating unserved requests.

The ratio of incoming requests to fleet size is the most influential variable on the performance of an aTaxi fleet. Because this ratio was kept constant in the simulated scenarios (as good as possible), the difference between the scenarios is rather small. The hourly demand profile is by far more dominant (Fig. 4). However, during the traffic and demand peak hours, a difference between the scenarios can be detected. During these hours, the travel times are most sensitive on additional vehicle flows and therefore the scenarios with less traffic tend to smaller effective waiting times.

The effective waiting time in the peak hours is rather large. The factor 10 between trips and fleet vehicles is probably chosen too high for a satisfying customer service. However, the important point for this comparison of scenarios was the consistent scaling of demand and fleet size.

VI. COMPARISON WITH OTHER TRAFFIC MODELS

We elaborate on traffic implications on the fleet performance by comparing the microsimulation model with two different traffic representations. The fleet operation algorithms are the same in all cases. The average link-level travel times of the last 5 minute interval are used to compute the vehicle-customer assignments and estimated waiting times in all traffic models. The difference between the models is the simulation of vehicle movements.

In the microsimulation, each vehicle has its own travel time according to speed limits, other headway vehicles, and – probably most importantly in the inner city – traffic lights and turns. Therefore, the estimated and the measured waiting times are distinct and can differ by a large amount, e.g. if there is a large time span between the estimation and the actual driving, if there are sudden changes in the traffic network, or if there are many turns on the route. The delay times for different turns at intersections are not estimated in the current model, because the link-level statistics are averaged over all lanes.

These effects distinguish the microsimulation from other traffic models used to simulate aTaxi operations. In the following section, we will describe two simpler traffic models and then highlight some consequences of the different traffic representations on fleet operations.

A. Description of Other Models

From a continuous map with constant travel time in [2], over a gridded network with 4 zones and peak / off-peak link-level travel times in [11] and real street networks with time-constant link-level travel times in [10] to real street networks with time-dependent link-level travel times in [5], there are many variations of traffic representation used in aTaxi simulations. There are also a few papers using flow-dependent traffic models as described in Section II.

We decided to compare our approach to the model of [10] and [5], namely a model with time-constant link-level travel times and a model with time-dependent link-level travel times. This was very convenient because the time-dependent network statistics of the base-simulation have already been available and the network already carries the information of the speed limits.

As in [10], the time-constant travel time was chosen to be 4/3 of the free-flow travel time, i.e. vehicles drive with 75% of the speed limit.

The fleet operation algorithm uses current link-level travel times for route computations in both cases. In the time-constant case, the vehicles drive that long as well. In the time-dependent case however, the driven travel times can be different if the traffic conditions changed in the time between the estimation and the time the vehicle actually drives.

We refer to the models by 'NoMicroSim_75pc' for the time-constant version, 'NoMicroSim_TT' for the time-dependent version, and 'MicroSim' for the simulation with the traffic microsimulation from now on.

B. Results

Fig. 5 displays the three defined types of waiting times (estimated, measured, effective) averaged over the whole simulation period for the different scenarios and simulation

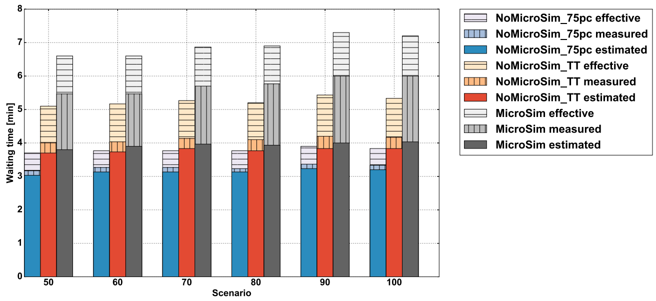

Fig. 5. Time-averaged customer waiting times in the different models.

models. Hereby, estimated and measured waiting times are only displayed for customers which are served. Therefore, the difference between measured and effective waiting time is directly proportional to the number of unserved requests.

Besides a constant ratio of fleet size and demand, the travel times are the same in all scenarios of the 'NoMicroSim_75pc' simulations. Thus, it is not surprising that waiting times are insensitive to the scenario. The estimated and measured waiting times are practically the same, while the rather small difference to the effective waiting times originates from the small amount of unserved requests.

The simulations using 'NoMicroSim_75pc' underestimate the travel times during peak hours compared to the more realistic travel times from 'NoMicroSim_TT'. In the latter, there is a small difference between estimated and measured waiting times because the estimation is made before the actual trip and travel times tend to increase with time between 5 and 11 am. Furthermore, the amount of unserved requests increases as well because there are more estimated waiting times above the threshold.

The estimated waiting times and penalty times due to unserved requests are nearly the same for the 'NoMicroSim_TT' and 'MicroSim' computations. This is consistent with the results from Section V, as traffic conditions hardly change due to the aTaxi operation. However, one would expect lower estimated waiting times for the scenarios where traffic conditions are better (50-90). The reason for the higher expected waiting times are the delays from previous jobs, which are reported to the fleet operation. The updated availability of vehicles is lower and causes slightly higher expected waiting times.

Delays are a realistic phenomenon. They can be caused by traffic incidents, customers not leaving the vehicle within the assumed boarding time, or something trivial like a fleet vehicle arriving just late for a green phase of an intersection. These delays are depicted by the large disparity between estimated and measured waiting times. The authors identified the following reasons for the deviations:

1) Although general delays at intersections can be captured by link-level travel times, the influence of different turns can not be represented.
2) Customer-vehicle assignments are fixed up to 10 minutes in advance; a delay-management with re-assignments would be able to reduce the disparity. Reducing the

maximal allowed waiting time would also reduce the underestimation of waiting times.

3) In this model, the vehicles are not routed according to the route computation that was used to generate the waiting time estimation.

VII. Conclusion

We illustrated the implementation of an autonomous taxi system into an existing, calibrated traffic microsimulation model. This approach was used to determine the impacts of an aTaxi system on the traffic network for selected scenarios in the city of Munich. Even in the scenario, where 10% of the the private vehicle trips originating and ending in the study area were exclusively replaced by aTaxi requests, the extra load due to empty rides increased network-wide private vehicle delay times by 1% only. This result is in line with Rossi et al. [8], who stated that rebalancing should not affect the network performance in a seriously bad way. However, the study by Maciejewski and Bischoff [9] displayed large delays being induced for higher replacement rates. We plan to test higher replacement rates in the future to investigate this discrepancy in our simulation environment.

Until autonomous taxis become a reality, the demand, and thereby the impacts on the traffic system, can only be estimated. This work investigated a few scenarios out of a large parameter space: The share of private vehicles could be even smaller. In [9], replacement rates started from 20% and went up to 100%. Moreover, there is the possibility that the extra comfort of autonomous vehicles attracts previous public transportation users and induces more traffic on the streets, be it in form of private vehicles or autonomous taxi rides. Of course, this would have negative consequences on the street network, and should not be the aim of autonomous taxis.

Furthermore, we estimated the effect of people changing their mobility behavior when using the aTaxi system. Most people, who abolish their private vehicles, will not use the aTaxis exclusively, but will hopefully travel with public transportation as well. We modeled this by splitting 10% of replaced private vehicle trips originating and ending in the study area to aTaxi and public transportation (by removing the request from the aTaxi system). The average delay times already decreased in the scenario where 1 out of 10 requests was handled by public transportation. If more than 2 of 10 requests used public transportation, the average delay time in the network decreased by more than 2%. This is a rather large value considering that we only modified within-area trips and left commuter and transit traffic unchanged, which account for more than 50% of the mileage in the network.

Additionally, we examined the consequences of different traffic models on the fleet operation. As expected, a model using time-constant link-level travel times based on free flow velocity (75% of it to be exact) overestimates the fleet performance in peak hours. Interestingly, the usage of a microsimulation instead of time-dependent link-level travel times has a larger effect than the additional traffic volume induced by empty rides. Waiting times vary slightly in the scenarios, which indicates that the induced traffic has only little influence. However, there are large delays in customer waiting times measured in the microsimulation compared to estimated waiting times computed by the operator based on link-level travel times. The process of averaging over all lanes (especially over all turn possibilities in front of intersections) and all signal phases disguises longer delays, e.g. for a left turn at a large intersection.

Moreover, we plan to improve the route computations by gathering turn-specific travel times from the aTaxi fleet.

Because of the shorter computation time, simulations based on link-level travel times are more suited in most cases, especially if lots of scenarios are to be examined, e.g. if fleet operation parameters or strategies are to be tested. The saved computation time is approximately given by the time of the base microsimulation model without any aTaxis. However, we want to highlight the use-cases for the presented framework of an agent-based simulation of an aTaxi operation in a traffic microsimulation:

- City planners: Studying traffic implications of aTaxi systems with different aTaxi penetration rates on their street network.
- Fleet operators: Testing fleet operation strategies under more realistic traffic conditions and checking for delay-robustness.

Acknowledgment

The authors would like to thank the Federal Ministry for the Environment, Nature Conservation, Building and Nuclear Safety for funding the project "City2Share".

References

[1] E. Martin and S. Shaheen, "Impacts of car2go on Vehicle Ownership, Modal Shift, Vehicle Miles Traveled, and Greenhouse Gas Emissions: An Analysis of Five North American Cities", *Transportation Sustainability Research Center (TSRC), UC Berkeley*, 2016.

[2] L.D. Burns, et al., "Transforming personal mobility", *The Earth Institute, Columbia University*, 2013.

[3] D.J. Fagnant, et al., "Operations of Shared Autonomous Vehicle Fleet for Austin, Texas, Market", *Transportation Research Record*, issue 2536, pp. 98-106, 2015.

[4] K. Spieser, et al., "Shared-Vehicle Mobility-on-Demand Systems: Fleet Operator's Guide to Rebalancing Empty Vehicles", *Proc. 95th Annual Meeting of the Transportation Research Board*, 2016.

[5] P.M. Boesch, et al., "Required Autonomous Vehicle Fleet Sizes to Serve Different Levels of Demand", *Proc. 95th Annual Meeting of the Transportation Research Board*, 2016.

[6] K. Spieser, et al., "Vehicle Routing for Shared-Mobility Systems with Time-Varying Demand", *IEEE American Control Conference (ACC)*, 2016, pp. 796-802.

[7] M.W. Levin, et al., "A General Framework for Modeling Shared Autonomous Vehicles", *Proc. 95th Annual Meeting of the Transportation Research Board*, 2016.

[8] F. Rossi, et al., "Autonomous Vehicle Routing in Congested Road Transportation Networks", *submitted to Autonomous Robots*, 2017.

[9] M. Maciejewski and J. Bischoff, "Congestion Effects Of Autonomous Taxi Fleets", *submitted to TRANSPORT*, 2016.

[10] W. Burghout, et al., "Impacts of Shared Autonomous Taxis in a Metropolitan Area", *Proc. 94th Annual Meeting of the Transportation Research Board*, 2015.

[11] T.D. Chen, et al., "Operations of a Shared, Autonomous Electric Vehicle Fleet: Implications of Vehicle & Charging Infrastructure Decisions", *Proc. 95th Annual Meeting of the Transportation Research Board*, 2016.