International Scientific Conference on Mobility and Transport Transforming Urban Mobility, mobil.TUM 2016, 6-7 June 2016, Munich, Germany

Autonomous taxicabs in Berlin – a spatiotemporal analysis of service performance

Joschka Bischoff[a]*, Michal Maciejewski[a,b]

[a]Department of Transport Systems Planning and Transport Telematics, TU Berlin, Salzufer 17-19, 10587 Berlin, Germany
[b]Division of Transport Systems, Poznan University of Technology, Piotrowo 3, 60-965 Poznan, Poland

Abstract

Autonomous taxi (AT) fleets have the potential to take over a significant amount of traffic handled nowadays by conventionally driven vehicles (CDV). In this paper, we simulate a city-wide replacement of private cars with AT fleets of various sizes.
The simulation model comprises microscopic demand for all private car trips in Berlin (including incoming and outgoing traffic), out of which the internal ones are exclusively served by ATs. With a fleet of 100,000 vehicles the city will be served at an appropriate level of service. Waiting times for an AT will generally be higher in the outskirts where the distance in between rides is higher. Additional traffic originating from ATs driving empty to pick up the next customer is here also expected to be considerably higher. Further results also suggest that a reduction of the service area to cover only the city center might be more efficient and that additional demand for AT services by people switching from public transit needs to be met by a proportionally bigger fleet size.
© 2016 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license
(http://creativecommons.org/licenses/by-nc-nd/4.0/).
Peer-review under responsibility of the organizing committee of mobil.TUM 2016.

Keywords: Autonomous vehicles; autonomous taxis; MATSim; fleet size; large-scale taxi simulation

1. Introduction

The market entrance of fully autonomous vehicles (AVs) might change currently established car ownership models substantially within the next decades. In the future, people may prefer to use autonomous taxicabs (ATs) over owning a private car to satisfy their urban mobility needs. On one hand, this would lead to a reduction in the

* Corresponding author. Tel.: +49-30-3142952; fax: +49-30- 31426269.
  E-mail address: bischoff@vsp.tu-berlin.de

total number of vehicles and the required parking space in cities, as only one big fleet of ATs would be required. On the other hand, empty vehicles driving between dropping off a customer to the next customer's pick up point might lead to an overall increase of traffic volume. It might also question the existence of certain forms of public transportation, such as busses, in less dense populated areas and lessen the attractiveness of public transport in general.

This paper aims to determine in which city regions AT services should be provided, how big the influence of a shift of public transit users towards AT services would be and what influence empty driving AT vehicles have on the overall traffic in the city.

2. Related Work

Current Research in AV development is both driven by the classical automobile industry as well as companies originating in the IT sector. Most of them are currently focusing on developing conventional cars with some autonomous functionality, such as adaptive cruise control or lane keeping. There is, however, a tendency to offer AT services in lieu of selling conventionally driven vehicles (CDV) in the longer term. Google's self-driving car is specially designed towards offering such services (Lee 2015), Toyota plans on operating an AT fleet as early as 2020 (Demetriou 2015), and most recently General Motors and Lyft have teamed up to cooperate in AT services (Hsu 2016).

The impact of such services on inner-city mobility has so far been discussed mainly in terms of fleet sizes and environmental aspects. Depending on the constraints set and the location used, most literature suggests that demand for CDV trips can be handled with AT fleets of 9 to 15% of the car fleet size, with more densely populated areas requiring less vehicles (Burghout et al, 2015) and more wide-spread cities requiring more (Burns et al, 2013). For Berlin, the authors have previously conducted a large-scale simulation study, concluding that 100,000 ATs will be sufficient to replace all inner city car-trips (Bischoff & Maciejewski, 2016). While some of the study's outcome has been used as the input for this paper, the case studies and questions researched in this paper are, to our knowledge, new and have not been conducted using simulation of such a high level of detail and large scale anywhere else so far.

3. Methodology

In this research, the open source transport simulation MATSim (Horni et al., 2016) and its Dynamic Vehicle Routing Problem (DVRP) extension (Maciejewski, 2016) have been used for simulation. Being an agent-based simulation, MATSim's mesoscopic traffic flow model allows a sufficiently detailed detail of simulation at high simulation speed.

3.1. MATSim and DVRP

MATSim is co-developed by TU Berlin and ETH Zürich. The basic concept of MATSim is the simulation of people (agents) along their daily routines (plans). These consist of activities (such as home or work) and travel (legs) in between them. In combination with scoring and altering plans using co-evolutionary algorithms, this three-step process is applied to a synthetic population. After several iterations some form of equilibrium is reached. Simulation of traffic flow is based on a queue model. Typically, car and public transport are simulated physically in a network. Depending on the case study, a fraction (such as 1, 10 or 25%) of the real-world population may be enough to get meaningful simulation results, with network capacities being adjusted accordingly.

Simulation of on demand transport services in MATSim is handled by the DVRP extension. ATs are coordinated by a dispatching service that reacts to incoming events (such as new request submissions, vehicle arrivals and departures) and dynamically re-optimizes ATs' routes and schedules in order to ensure efficient processing of taxi demand. The dispatching strategy used in this study either assigns the closest taxi to a customer (in times of oversupply) or the closest customer to an empty vehicle (in times of undersupply, e.g. during peaks). It has been thoroughly described in previous studies (Maciejewski & Bischoff, 2015, Maciejewski et al., 2016). DVRP simulations in MATSim are usually carried using the whole sample size of the population, rather than just a fraction.

While this does not cause any computational problems with taxi fleets of today's dimensions, some adaptions were needed to serve millions of customers and dispatch hundreds of thousands of vehicles. These enhanced dispatch algorithms are described in (Bischoff & Maciejewski, 2016).

3.2. Initial model and scenario adaptation

The demand for AT trips used in this paper is derived from the MATSim Berlin scenario originally described by Neumann (2014). It depicts the synthetic population of a typical weekday in Berlin based on survey data from 2008. Car traffic in the scenario is characterized by two peaks, one smaller in the morning and one bigger in the afternoon. The modal split between car and public transit trips is roughly even. The scenario comes in different sample sizes, with the 10% scenario being the most commonly used (Kaddoura, 2015).

AT trips are initially derived from the 100% scenario by switching all car trips starting and ending within the city borders and removing all trips made by public transit and all those car trips that do not touch the city boundaries. This results in roughly 4.7 million trips, of which 2.5 million are made using AT and the rest by car. This setup, as used previously (Bischoff & Maciejewski, 2016) is referred to as the *100% base case*. For this case, 100,000 AT are sufficient to serve the city's transport demand with average wait times of less than three minutes. On average, each vehicle spends around 7.5 hours a day serving customers. Further statistics about this base case may be found in the corresponding study.

Since computational times for the 100% base case are relatively high (above 3 hours on a computer with the Intel Core i7-3930K processor) and the huge amount of output data being rather unhandy to process, we also re-created the base case using the original 10% scenario. However, the amount of trips does not exactly scale down. With 278,000 AT trips, the 10% scenario consists in fact of 11% of demand of the 100% base case. For naming conventions, the name *10% base* case is kept nevertheless.

In general, it is unclear whether a direct downscaling of fleet sizes results in similar simulation results as the 100% case. While usually MATSim scenarios can be downscaled without significant side effects, it has never been done in scenarios where forms of dynamic transport are used. With the overall number of trips in the 10% Berlin scenario remaining rather high, we assumed (and later confirmed) downscaling should work. Nonetheless, a comparison of both base cases is conducted in Section 4.1 to test the approach.

The simulation of background traffic (incoming and outgoing traffic) is handled by adjusting network link speeds dynamically over time according to the simulation results obtained for the original scenario. The direct combination of ordinary cars and public transit vehicles moving together in traffic requires assumptions regarding the co-existence of AV and CDV fleets that would go beyond the scope of this paper. Also not taken into account is the ongoing growth of population in Berlin, which, by the time ATs eventually could become available, might have an influence on results.

3.3. Relevant service criteria

Service criteria for autonomous taxicab operations are not necessarily any different to those of taxi services nowadays, besides the obvious irrelevance of driver related aspects and a more regular maintenance and cleaning interval. Criteria are generally set by both the customer and operator.

From a customer's perspective, waiting times for an AT should never be considerably higher than it takes to usually park and un-park a vehicle. Apart from the average waiting times, the 95 percentile of wait time is set as a criterion.

On the operator's side, minimizing fleet costs is of major importance. Besides the actual fleet size, time or kilometers spent driving empty must remain low. Furthermore, depending on the fleet concept, the operator may or may not want the driving load distributed evenly onto all vehicles.

3.4. Service assumptions

All AT requests are handed in by customers immediately before departure, there are no advance bookings. This not only reflects the most typical behavior in today's taxi market, but also corresponds to other forms of demand-responsive transport, such as free-floating car sharing, where longer pre-bookings are not generally accepted. Ingress and egress times are set to one and two minutes respectively. After delivering a customer, ATs are parked at the drop off position until the next dispatch. No ranks are used. AT service is available during the whole day. The initial fleet distribution in the morning is in accordance with the population density in each of the city's 447 statistical units (*LOR – lebensweltlich orientierte Räume (Senatsverwaltung für Stadtentwicklung und Umwelt, 2016)*).

4. Base case

In this section, the 10% base case (corresponding to 11% of the AT demand) is derived and analyzed in terms of fleet usage and the spatial distribution of empty rides and customer waiting times throughout the day.

4.1. Scenario downscaling

For the *100% base case* described in Bischoff & Maciejewski (2016), a fleet of 100,000 vehicles is sufficient to replace all inner city car trips. With this number of vehicles, the average wait time for a vehicle can be kept under three minutes at most times of the day. During peak hours it climbs to just under five minutes (cf. to Fig. 1). The 95 percentile of wait time is just above 14 minutes. As a direct comparison with the 10% base case in fig 2 suggests, the passenger wait times over the day are fairly similar for a fleet size of 11,000 (i.e. 11%, as noted in Section 3.2). In terms of passenger service, scaling down the scenario does not lead to a loss of information.

Fleet occupancy statistics over the day behave also quite similar (Figs. 3 and 4): The temporal distribution of idle vehicles (black dashed line), empty driving (blue line) and occupied vehicles (green line) follows the same patterns during the day in both cases. The red line, which displays requests that upon placing remain initially unplanned, shows that there is an undersupply of vehicles in the afternoon peak in both scenarios which is compensated by switching to a different dispatch mode (as described in Section 3.1). In both scenarios, this shortage of vehicles is overcome roughly equally fast. Duration of pickup trips and rides with customer are also of comparable (less than 3% difference) length.

Overall, the *10% base case* is quite comparable to the full scale model in terms of simulation outcome. Using it in further analysis is therefore not expected to have major impact on the validity of results.

4.2. Fleet usage

On average, each vehicle is driving for 274 km during the course of the day (median 272 km), of these, 239 km are with a passenger on board, resulting in an extra mileage of 35km (13%) spent driving empty. The average trip distance is 9.1 km. In terms of time spent driving, each vehicle is occupied for 6.8 hours per day, with 7.6 hours spent overall on the road. With a standard derivation of 88 km, some vehicles are significantly busier than others. Since no active balancing between vehicles is conducted and there are no constraints regarding vehicle range, this does not affect service.
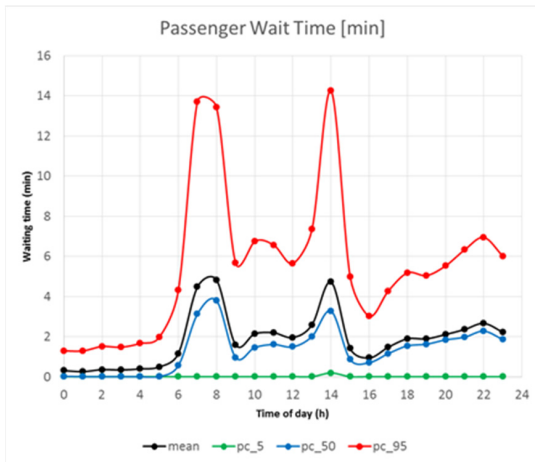
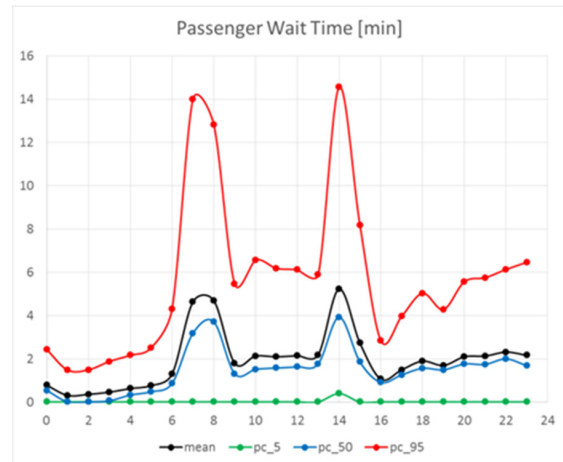Fig. 1 Average passenger waiting times for 100,000 vehicles in the 100% base case

Fig. 2 Average passenger waiting times for 11,000 vehicles in the 10% base case (≈11% demand)
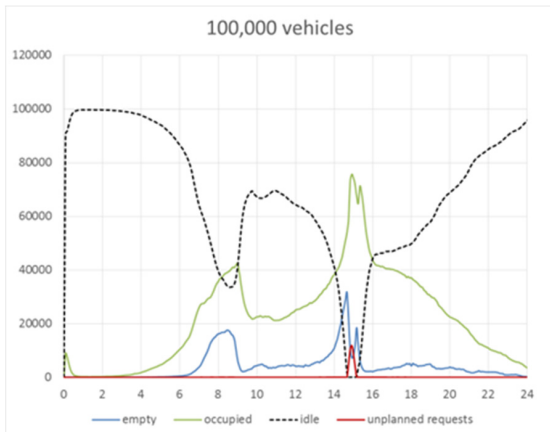

Fig. 3 Fleet occupancy and unplanned requests during the day in the 100% base case for a fleet of 100,000 vehicles
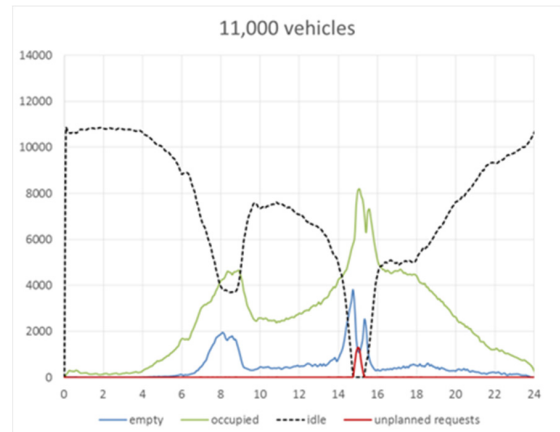

Fig. 4 Fleet occupancy and unplanned requests during the day in the 10% base case (≈11% demand) for a fleet of 11,000 vehicles

4.3. Spatial distribution of waiting times

Even during peak times, average waiting times seem acceptable. A closer look into the spatial distribution shows however a huge distribution of these across the city, as Figs. 5 and 6 reveal. In densely populated areas, waiting times are generally well below the hourly average both in the morning and afternoon. Especially during the morning peak, the relatively high demand from the outskirts leads to longer average waiting times in these regions (and thus also to longer empty rides). Some regions have average waiting times of almost 20 minutes. However, as displayed by the 95 percentile in Fig. 2 at 7 am suggests, less than five percent of the trips have more than 14 minutes waiting time, i.e. the overall share of trips from very remote regions is low. In the afternoon, the effect is somewhat milder, but, with average waiting times of up to ten minutes in some outskirts, still noticeable. However, the overall demand of trips originating from these zones is somewhat lower during these times than in the morning. In the city center, waiting times remain stable during the whole day.

4.4. Spatial distribution of empty rides

Empty rides will have a significant impact on traffic in the city. Even though in total only 16% of the overall drive time is spent driving without a customer, the additional mileage on the road network is not spread evenly throughout the city. As Fig. 7 shows, traffic generated by empty rides is generally expected to be less than 10% in the city center. This additional amount of traffic might have an impact on travel flow, but can most likely be handled by the higher efficiency AV operations will bring by taking advantage of better communication technologies between cars. In the outskirts, effects of empty driving vehicles will be higher in many areas. While congestion is not usually a bigger issue in these areas today, an empty ride share of up to 45% will most likely lead to negative effects here. Congestions caused by these additional trips might not have a direct impact on people riding taxis, but could have a negative influence on the living quality in these areas. This goes in line with the fact, that average pickup distances are also considerably higher in the outskirts (see Fig. 8): While this value is generally below or around 1 km in most city center areas, up to 5 km occur in less densely populated zones. These figures and numbers overall create the impression, that AT services will be of lower efficiency in sparsely populated areas.
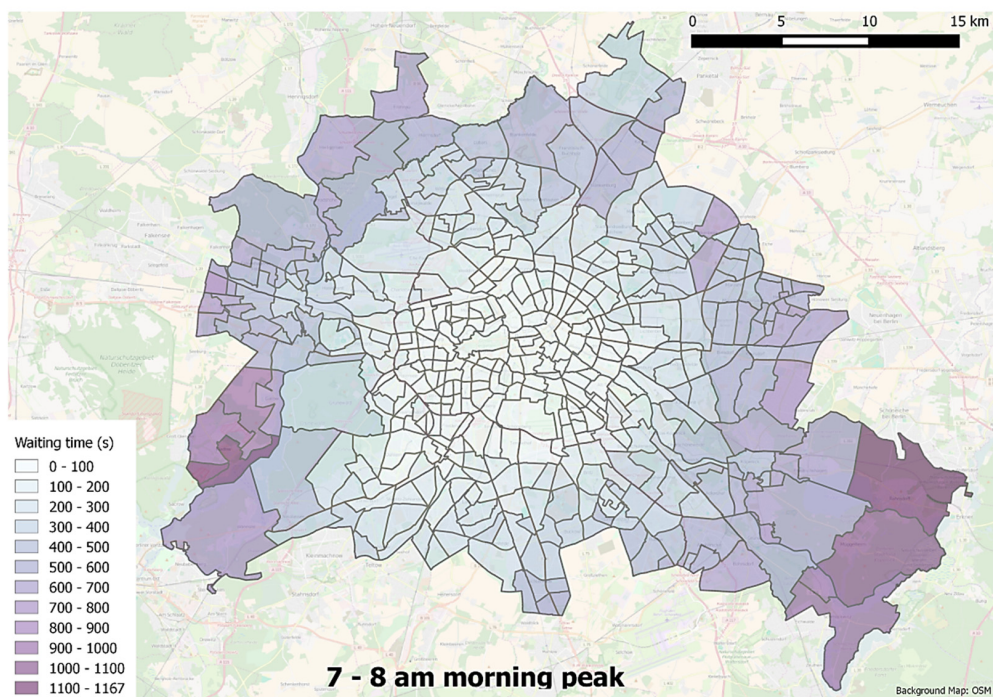
Fig. 5 Average waiting time per zone during the morning peak. Average waiting over all zones is just under five minutes.
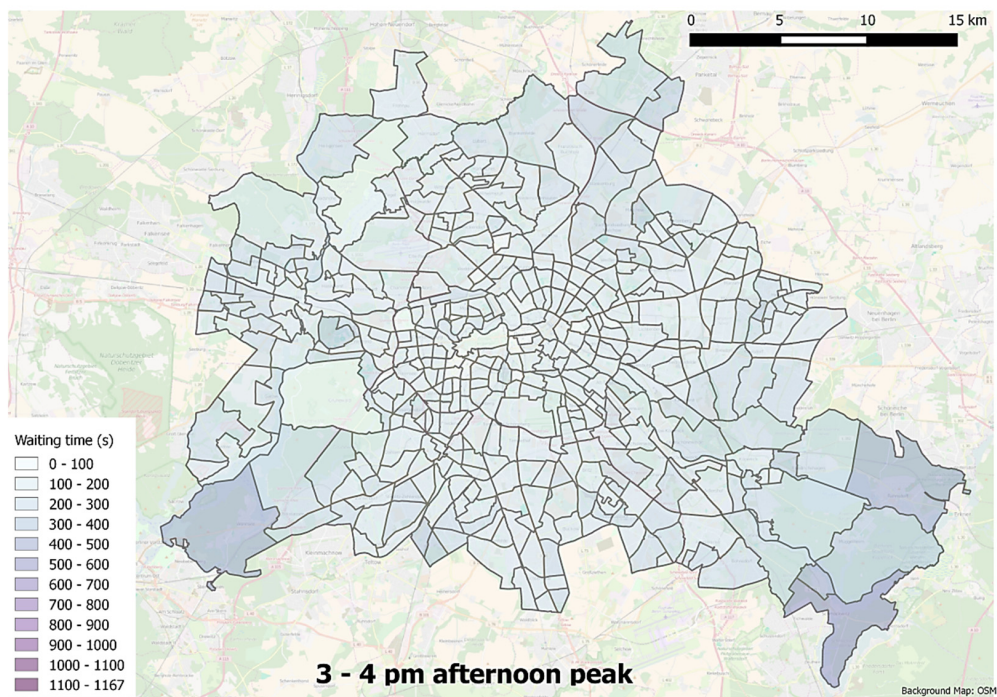

Fig. 6 Average waiting time per zone during the afternoon peak. Average time over all zones is around five minutes.
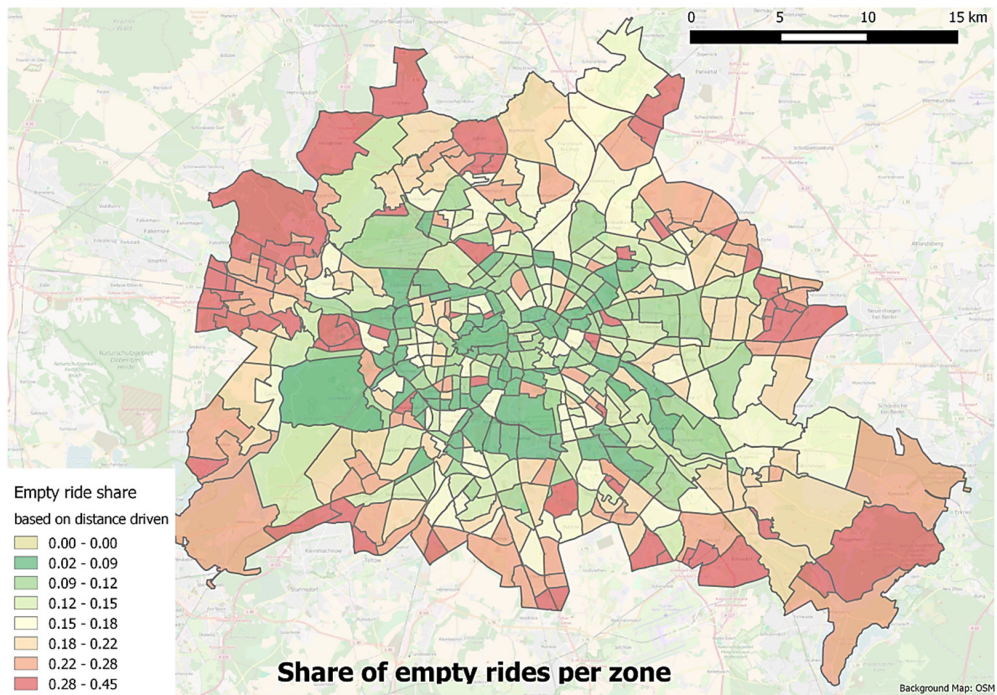
Fig. 7 Share of empty mileage driven empty per zone. The outskirts generally have a higher share of empty rides.
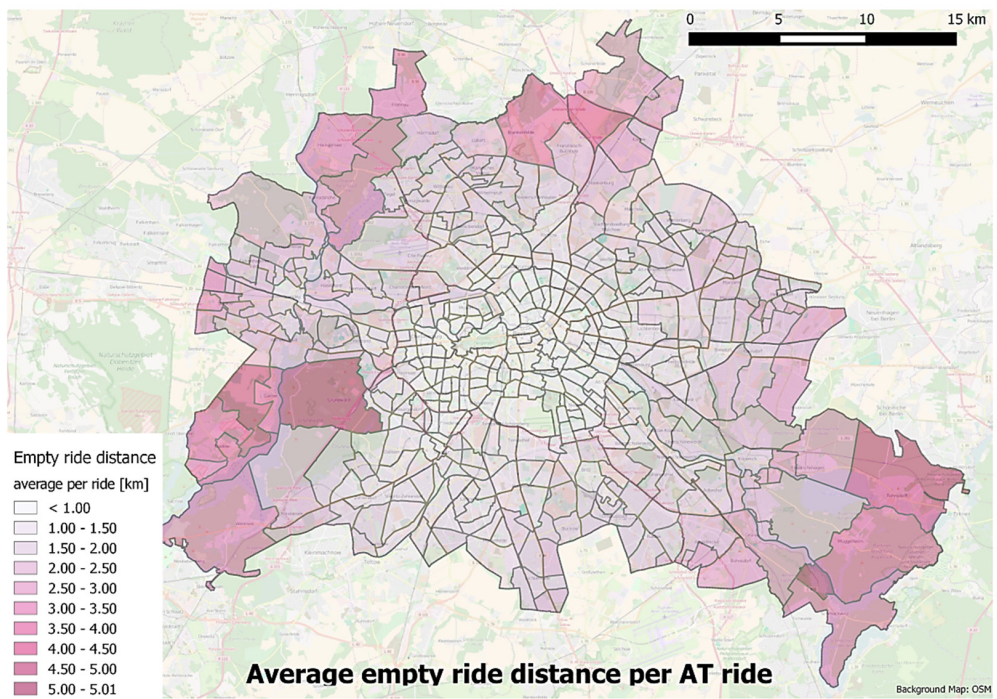

Fig. 8 Average empty ride distance per AT trip per zone. These are generally very short in densely populated areas.

5. Scenario variations

In this section, the base case is altered in two ways: Firstly, the area of service is reduced to the city center, and secondly, a certain percentage of the original public transit demand is shifted to AT usage.

5.1. A smaller service area and mobility hubs

As the results presented in Sections 4.3 and 4.4 suggest, the AT service is able to operate much more efficient within the city center than in the outskirts of the city. This raises the question if a smaller service area may not be a better choice for both the operator and overall city traffic. Hence the service area was reduced to the city center only and consists now roughly of the *S-Bahn* (urban rail) circle, an area that is often used for different kinds of transport-related policy cases in Berlin. This area is designated to be used by ATs only, meaning that passengers travelling into or out of the zone by car would requested to change from car to AT (or vice versa) at one of four mobility hubs near the border of the new service area. These hubs are located at trunk roads going into and out of the city center. Fig. 9 provides an overview of the setup. Overall, some 174,000 trips are dispatched in this scenario, of which 129,000 either begin or end at one of the hub locations. Car trips, that now take place only outside the AT service area, were not simulated.

To evaluate the fleet size required, several simulation runs with fleets between 4,000 and 8,000 vehicles were conducted. Results suggest that a fleet of 6,500 to 7,000 ATs would be necessary to cope with demand in a similar way as in the original scenario (i.e. with an average waiting time between two and three minutes during most times of the day). With 6,500 vehicles, each AT is handling 26.8 trips during the day (+1.7 compared to the base case). These trips are considerably shorter, with the average distance of almost 7 km. The resulting waiting time distribution is all in all rather constant throughout the city center (Fig. 10). However, the waiting times at hub locations are above average, which indicates that these need to be taken into account when redistributing empty ATs. Pickup distances are similarly low as in the base case for most zones, with some exceptions at the hub locations, stressing the question of initial vehicle allocation.

All in all, reducing the service area does increase the efficiency of the system. What remains debatable though is the question, in how far such a scenario with hubs and the need to change from car to AT in the middle of a trip will be realistic. While this scenario might require less street side parking spaces in the center, huge parking spaces at the hub locations will need to be created. It might be better to leave car traffic going into and out of these zones untouched and decrease the AV fleet accordingly.
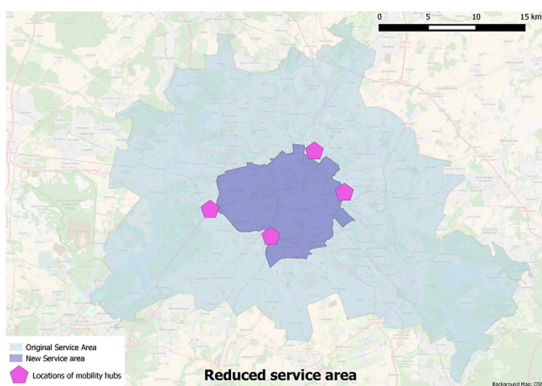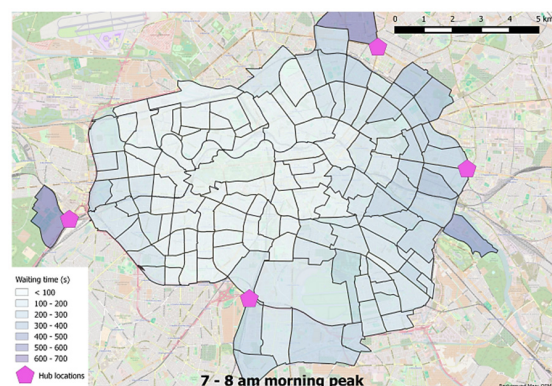

Fig. 9 Area of service and hub locations

Fig. 10 Waiting times during the morning peak for the hub case

5.2. Public transport shift

In the original scenario, the share of rides between the public transport and car modes is roughly even for rides within Berlin, as is their general structure (distance and timely distribution). Since AT services are available also to people for whom car ownership is not an option, a certain mode shift from public transport towards them can be expected. In this specific case, 10% of all public transit rides, independent of their original mode (train, subway, bus or tram), were transformed into AT trips.

This results in 323,000 trips made by AT. The fleet size was initially left at 11,000 ATs. Average all day waiting time for trips climbs from 2:35 min to 2:52 min. To compensate this effect, a fleet increase to 12,000 vehicles would be necessary (resulting in 2:39 average waiting time). An increased demand seems therefore compensable with a proportional increase in fleet size.

6. Conclusion

A large scale introduction of AT services in Berlin will require roughly 100,000 vehicles to serve the demand currently served by roughly 1.1 million cars. In this paper, we were able to show that for many purposes, a simulation using a 10% sample of fleet size and demand is providing sufficiently good results at very high computational speeds.

The simulation outcome also suggest AT services will be working most efficiently in densely populated city areas, where the influence of empty driving vehicles is expected to be rather low. In the outskirts, their influence will be much higher, with empty driving ATs making up for almost half of the share of AT trips in some areas. Maintaining private car ownership models might be preferable in such areas.

Decreasing the service area to cover only the city center will be of higher efficiency from a fleet operator's side. The combination with mobility hubs might form a possible option, but it seems debatable whether an introduction of a car-free zone where only ATs are operative forms a realistic scenario. Further research could focus on cooperation and integration effects of CDVs and ATs in the city center without the usage of hubs. This will require some more profound data on the effects of autonomous vehicles on traffic dynamics and, possibly, higher resolution simulation.

A very general shift from public transit towards ATs shows that fleet size will need to grow proportionally to demand increase. Further research should focus here on a more detailed study which forms of public transit and which lines are most likely to lose customers by introducing AT services.

Acknowledgements

References

Bischoff, J., Maciejewski, M.. Simulation of city-wide replacement of private cars with autonomous taxis in Berlin. Procedia Computer Science, 2016, 83, 237 - 244

Burghout, W., Rigole, P.J., Andreasson, I.. Impacts of shared autonomous taxis in a Metropolitan area. Transportation Research Board 94th Annual Meeting; 15-4000. 2015,

Burns, L.D., Scarborough, B.A.. Transforming personal mobility. Tech. Rep.; The Earth Institute of Columbia University; 2013.

Demetriou, D.. Japans self-driving taxis gearing up for 2020 Tokyo Olympics. The Telegraph,

02 Oct 2015; 2015. URL: http://www.telegraph.co.uk/news/worldnews/asia/japan/11906275/Japans-self-driving-taxis-gearing-up-for-2020-Tokyo-Olympics.html.

Horni, A.; Nagel, K.; Axhausen, K.W. (eds.) 2016. The Multi-Agent Transport Simulation MATSim, Ubiquity, London. Available online: http://matsim.org/the-book

Hsu, J.. GM and Lyft team up for robot taxi service. IEEE Spectrum, 04.01.2016; 2016. URL: http://spectrum.ieee.org/cars-that-think/transportation/self-driving/gm-and-lyft-team-up-for-robot-taxi-service.

Kaddoura, I.. Marginal congestion cost pricing in a multi-agent simulation: Investigation of the greater Berlin area. Journal of Transport Economics and Policy 2015;49(4):560–578

Lee, D.. Google's driverless car is brilliant but so boring. BBC, 02.10.2015; 2015. URL: http://www.bbc.com/news/technology-34423292.

Maciejewski, M., Bischoff, J.. Large-scale microscopic simulation of taxi services. Procedia Computer Science 2015; 52:358–364.

Maciejewski, M., Salanova, J.M., Bischoff, J., Estrada, M.. Large-scale microscopic simulation of taxi services. Berlin and Barcelona case studies. *Journal of Ambient Intelligence and Humanized Computing*, 2016, 7, 385-393.

Maciejewski, M., 2016. Dynamic Transport Services, in Horni et al. (Eds.), *The Multi-Agent Transport Simulation MATSim*, Ubiquity, London. Available online: http://matsim.org/the-book

Neumann, A.. A paratransit-inspired evolutionary process for public transit network design. Ph.D. thesis; Technische Universität Berlin; 2014.

Senatsverwaltung für Stadt und Umwelt: Lebensweltlich orientierte Räume in Berlin, 2016; URL: http://www.stadtentwicklung.berlin.de/planen/basisdaten_stadtentwicklung/lor/