Simulation-based analysis of the impacts of fleets of autonomous vehicles on urban traffic *

Kai Nagel, Joschka Bischoff, Gregor Leich, Michał Maciejewski

Fachgebiet Verkehrssystemplanung und Verkehrstelematik
Institut für Land- und Seeverkehr
TU Berlin

nagel|bischoff|leich|maciejewski@vsp.tu-berlin.de

February 18, 2019

1 Introduction

It is now conceivable that fully autonomous, i.e. self-driving vehicles may be admitted in urban traffic. Here we want to consider so-called "Level 5" vehicles, which will not even have a steering wheel. We will assume that these vehicles are not owned by individuals, but instead by one or more companies that operate them as fleets. Such vehicles are often called "robotaxis" or Shared Autonomous Vehicles (SAVs), where we will initially leave open how many passengers per vehicle we will allow.

SAVs may have a disruptive effect on both our transport system and on urban planning (Levinson & Krizek, 2015; Axhausen, 2016). Possibly, they will significantly increase the overall system performance of road-based transport. Compared to existing public transport services, they may offer a more passenger-oriented and cheaper service. In combination with Electric Vehicle technology, SAV fleets may operate without any direct carbon emissions. During the course of the last years, the authors have explored several questions related to SAVs using microscopic simulations. Rather than on behavioral questions of users,

*VSP Working Paper 19-06, see http://www.vsp.tu-berlin.de/publications/vspwp.

which remain hard to predict, given that no AV service is currently operative, these simulations primarily focus on the "mechanical" aspects of such fleets. Those may include waiting times for vehicles depending on fleet sizes, estimations about electric energy consumption and power demand at chargers and how feeder services to train stations may look like.

This paper provides an overview over these simulations. It starts with a section about the simulation technologies (section 2), followed by an introduction into simulating fleets of SAVs (section 3). Section 4 discusses different dispatch algorithms for these vehicles, followed by an investigation of the effects of a full replacement of conventional cars within the city boundaries of Berlin (section 5). Such a replacement will initially result in additional vehicle miles traveled (VMT) and thus to more congestion. In section 6 it is then discussed how an increased flow capacity of AVs will affect the roads. In section 7, the possibilities and effects of an electrified SAV system are discussed. The following two sections 8 and 9 focus on a pooling of passengers in vehicles. The paper concludes with a discussion (section 10) and closes with policy-related conclusions and related questions (section 11).

2 Microscopic simulation of urban traffic

For many years, we have now programmed and used microscopic simulations of traffic of cities and regions. 'Microscopic' means that each person, each vehicle, each intersection etc. is individually resolved by a synthetic avatar in the simulation. In particular the persons are synthetically generated objects, which represent the true population only in the statistical sense. Each synthetic person obtains at least one activity-based daily plan, for example "home – work – shop – leisure – home". Each activity obtains a location and a time at which the activity ends.[1] Activities at different locations need to be connected by trips, which will be undertaken by a mode of transport. In the software MATSim (Horni et al., 2016, Multi-Agent Transport Simulation) this is encoded approximately as follows:

```
<population>
  <person id="1">
```

[1] A simple option to generate such daily plans is from travel diary surveys, as for example the German SrV (Ahrens et al., 2014, System repräsentativer Verkehrsbefragungen) or MiD (Infas & DLR, 2010, Mobilität in Deutschland). In those trip diaries, the activity locations are normally given through relatively coarse zones; it is often sufficient to randomly select a point within the zone, where ideally the randomness is weighted by land use data.

```
    <plan>
      <act type="home" x="5.0" y="8.0" end_time="08:00:00" />
      <leg mode="car" />
      <act type="work" x="1500.0" y="890.0" end_time="17:30:00" />
      <leg mode="car" />
      <act type="home" x="5.0" y="8.0" />
    </plan>
  </person>
  <person id="2">
     ...
  </person>
</population>
```

MATSim generates the missing routes and then starts an iterative process that consists of the following steps:

1. Every synthetic person has one "selected" plan (and maybe other inactive plans, see below).

2. **MobSim:** A traffic flow simulation, also called MobSim = mobility simulation or synthetic reality, executes all selected plans simultaneously in the synthetic reality. This results, for example, in congestion or crowded public transport vehicles.

3. **Scoring:** Each synthetic person afterwards scores the plan's performance in the simulated day. This scoring often corresponds to an econometric utility function, and contains positive contributions for time spent at activities, and negative contributions for time spent traveling. If desired, one could specifically penalize time in congestion or in crowded public transport vehicles.

4. **Replanning:** All synthetic persons may revise their selected plans. Some of the persons generate a new plan, for example with another route, another mode of transport, other departure times, other activity locations, or another activity sequence. These newly generated plans are made "selected" in the sense of item 1; since the synthetic persons memorize all previously executed plans, this means that over time each synthetic person developes a choice set of multiple plans. Those persons that do not generate a new plan select between their memorized plans, typically with a logit model. (Flötteröd & Kickhöfer, 2016).

   This process is then repeated by returning to step 2.

The iterations are terminated when the synthetic persons find better alternatives only rarely.

3 Integration of SAVs

On the demand side, the integration of SAVs is straightforward: one simply sets either initially or in step 4 the mode of transport to "SAV". On the supply side, it is more complex: When in the traffic flow simulation, step 2, a synthetic person ends her activity to undertake a trip by SAV, then a request for such a vehicle is communicated, and the person waits for the vehicle's arrival. A dispatch center receives the request, selects a vehicle, and sends it to the customer. The SAV drives to the customer, lets her board, drives to the destination, and lets the customer alight. In the present implementation, the vehicle then stays at that location until contacted again by the dispatch center. Alternatives to this are discussed in section 10.2.

4 Dispatch

One central aspect for the efficient assignment of vehicles to customers is the actual dispatch algorithm. For single-ride requests, we have studied several algorithms over the years. These include:

Closest idle vehicle – For each incoming request, the idle vehicles are scanned, and the one that could be quickest with the customer is selected.

This on first sight plausible algorithm is, on second sight, problematic under high load: When there are no idle vehicles left, then the incoming requests will be queued, and processed in their sequence. Every time a vehicle becomes idle, it will be paired with the oldest unassigned request. Evidently, it can now easily happen that vehicle and request are far apart.

Closest vehicle overall – In contrast to the last approach, this one here includes busy vehicles. It computes between *all* vehicles which one would be fastest with the customer, including the remaining expected travel times with customers if any. Once assigned, vehicle dispatch remains fixed.

Opportunistic – Here the sort order is not by incoming request, but by vehicles becoming idle: Each vehicle that just became idle is assigned to the nearest open request. – It is evident that this approach reduces empty trips. It also may be unfair, as trip wishes originating from areas where few trips are currently ending, will incur longer waiting times.

Optimal assignment – Rather than the fixed assignments discussed so far, one can re-optimize the assignment frequently, e.g. every time when a

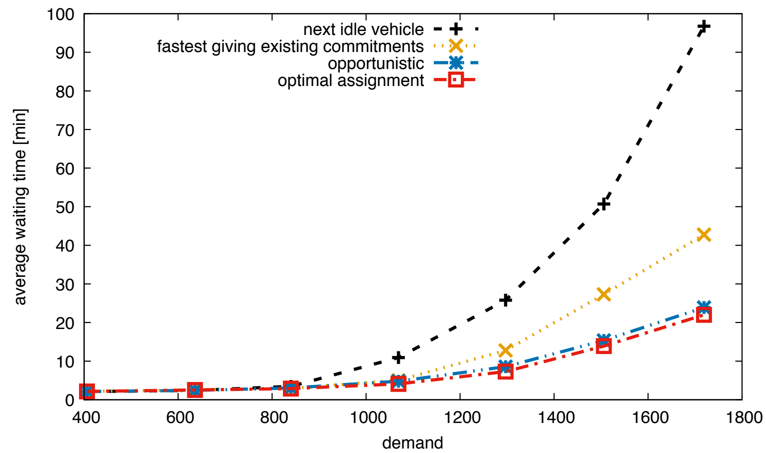

Figure 1: Average waiting time for a SAV in a simulation with different dispatch algorithms as a function of demand. With rising demand, simple heuristic dispatch algorithms result in longer waiting times. Maciejewski & Nagel (2013a,b); Maciejewski (2014a,b).

> request comes in, every time a vehicle becomes idle, and every time a vehicle is much slower or faster than assumed previously. One attractive approach here is the Hungarian method (Kuhn, 1955), where the cost between each request and reach vehicle is simply the (expected) time for the vehicle to reach the customer, possibly including the remaining time with the current customer. As is standard, the problem can be padded with dummy vehicles or requests when the numbers do not exactly match. The Hungarian method computes an optimal assignment, but the approach is still a heuristic since it does not look more than one step (trip) into the future.

One can study these approaches under different loads (Fig. 1). When the load is low, all of these approaches perform similarly; the empty drive time to a customer just depends on the density of the vehicles in the system. Under high load, as expected the approach "closest idle vehicle" quickly accumulates backlog and thus leads to long waiting times. "Closest vehicle overall" is better, but the opportunistic approach is again even better under high load. This is intuitively clear, since the opportunistic approach reduces empty drive times, thus maximizing system throughput, since the customer trips are the same with all approaches. The downside is that the opportunistic approach is potentially unfair, since requests are no longer served in the sequence of their arrival, and

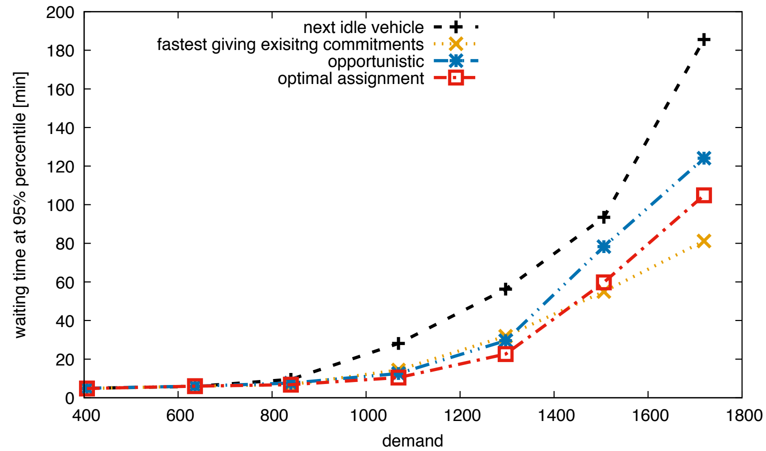

Figure 2: 95% percentile of highest waiting times. The opprtunistic approach is clearly dispatching vehicles rather unfairly, as five percent of customers have to endure long waiting times. Only the "closest idle vehicle" produces even worse results Maciejewski (2014a).

thus customers far away from any dropoffs will be left stranded until the peak period is over; see Fig. 2 for the waiting times at the 95% percentile. Optimal assignment is only slightly better, implying that what it effectively does is similar to the opportunistic approach.

It is plausible to attempt so-called off-line algorithms for comparison. Here, all requests are known beforehand. In theory, this should score very good results. However, in our simulations we could see that these often perform *worse* than the on-line approaches described above. This can be explained by the fact that travel times are stochastic. A pre-computation using off-line algorithms is not able to react to these fluctuations in travel times. Only if all travel times are also known in advance, an off-line algorithm is able to score significantly better results than the on-line solutions described here. We were, however, unable to compute optimal solutions within acceptable computing times (Maciejewski, 2014b). Overall, using on-line approaches seems therefore plausible, even if all requests are known in advance.

Based on the investigations described above, the following text will use the approach "closest vehicle overall" as long as all requests can be served, and the opportunistic approach under high load. We assume real world loading curves, which means that any peak period is eventually over and the backlog can be served. Alternatively, but currently not investigated by us, one could reject

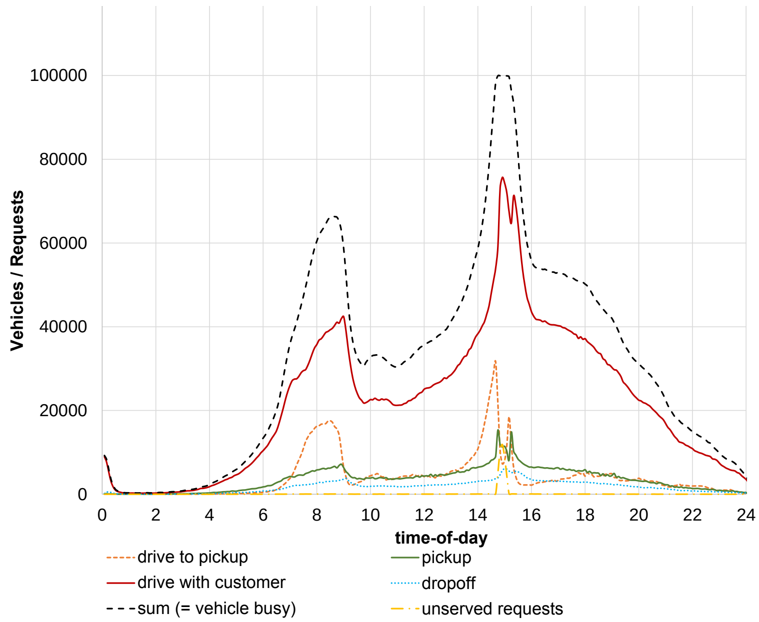

Figure 3: Vehicle states as function of time-of-day.

requests, or customers could re-plan when confronted with long waiting times.

5 Replacement of all private vehicles in Berlin by a fleet of autonomous vehicles

A question now is how an urban system with a fleet of autonomous vehicles would perform. We have, therefore, taken an existing model for the traffic of Berlin, and replaced all private car trips within the city boundaries by robotaxi trips (Bischoff & Maciejewski, 2016b). All trips to or from outside the city boundaries are still performed by private cars. We have performed simulations with different fleet sizes, and finally settled on a fleet of 100 000 vehicles, for a reason described below. We have also assumed that there are no additional mode switch effects.

Fig. 3 shows the different vehicle states as a function of the time-of-day. For each time, the number of vehicles in the different states "empty drive" (= on the way to customer), "pickup" (= passenger boarding), "dropoff" (= passenger alighting), and "occupied" (= trip towards destination with customer on board). The sum of these four numbers results in the number of busy vehicles; all other vehicles

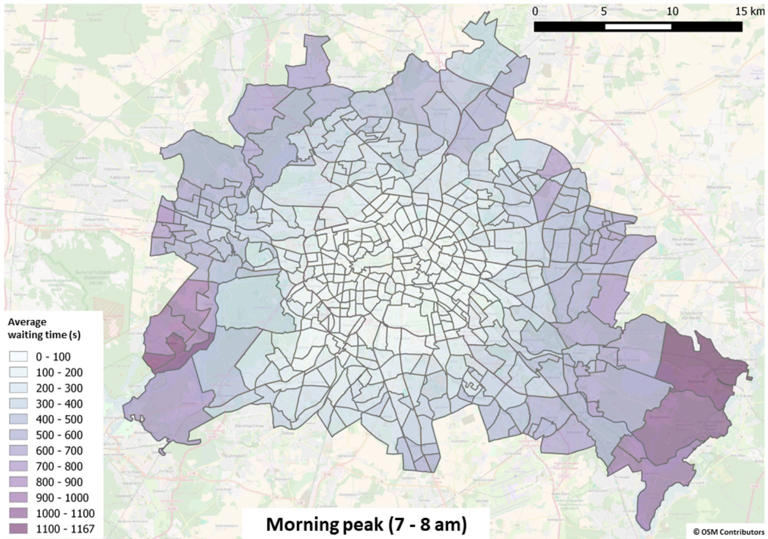

Figure 4: Spatial waiting times (morning). Cf. Bischoff & Maciejewski (2016a).

are idle. Additionally, the plot shows around 3pm the number of unassigned requests because of high load; in Berlin, the afternoon peak has higher load than the morning peak. From the load at this point in time we took the fleet size of 100 000: a fleet size that is just barely sufficient to handle the peak demand. One clearly observes that large parts of the fleet are idle during much of the day. This implies further potentials, e.g. for small-scale freight traffic.

The 100 000 fleet vehicles would replace about 1 million privately owned cars. The corresponding parking areas would no longer be needed for that purpose and could thus be made available for other usages.

Fig. 4 shows the resulting waiting time in Berlin between 7 and 8 am. One clearly recognizes increased waiting times of up to 15 minutes in the outer zones. A first wave of vehicles towards the workplaces in Berlin has already taken place here, so that most of the vehicles are in the inner city, and need time to travel back to the outer zones. In the afternoon, the effect is *not* the same (Fig. 5): because of generally many trips in the inner city, there is always a sufficient number of vehicles becoming idle, even during the afternoon peak.

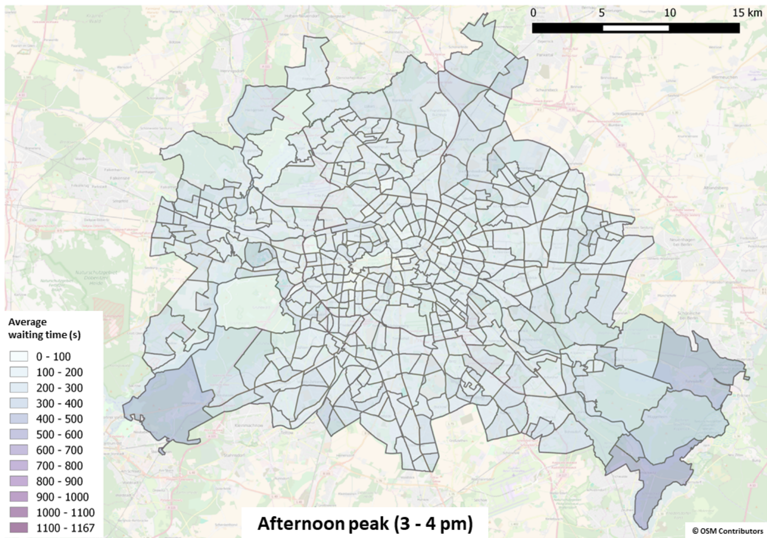

Figure 5: Spatial waiting times (afternoon). Cf. Bischoff & Maciejewski (2016a).

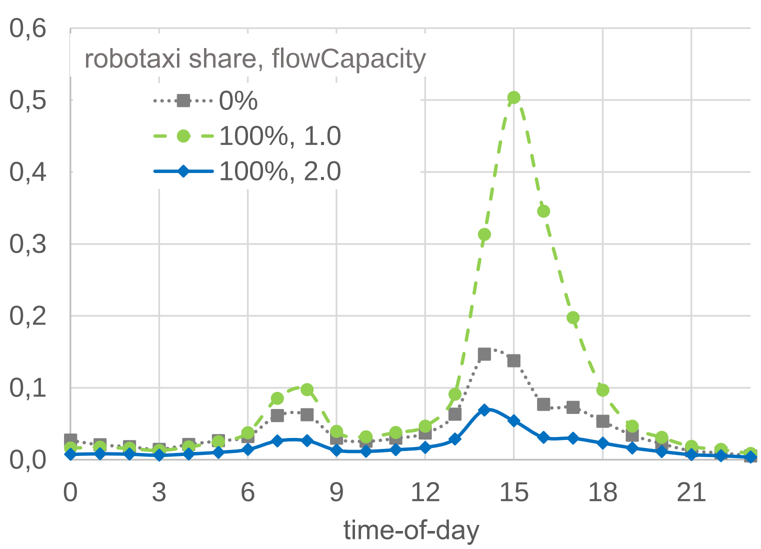

Figure 6: Relative travel time increase as function of time-of-day.

6 Congestion effects

The scenario of Sec. 5 generates additional vehicle kilometers, since the vehicles need to do the empty trips from dropoff to pickup in addition to the trips with customers. For the described simulation, we found this to be an additional 13% of vehicle kilometers (Bischoff & Maciejewski, 2016a, Sec. 4.2). This results in a strong increase of time losses due to congestion from 15% to 50% during the peak period (see Fig. 6). A number of 15% means that the trips in average need 15% more than in the empty system. This base case number is relatively small since the simulation includes a large part of the surroundings of Berlin.

In this situation, one can speculate that autonomous vehicles will consume less flow capacity than conventional vehicles. Evidently, the capacity gain needs to be at least 13% to compensate for the additional empty trips. Fig. 6 shows the consequences of a possible flow capacity improvement of a factor of two. Our simulations are not sufficiently detailed to generate the possible capacity gains by autonomous vehicles endogenously; we thus need to use results from other researchers. Still, we would like to mention that this may be as much a question of culture as of technology: For example, one could retrofit the existing conventional vehicles by relatively cheap RFI (radio frequency identification) transponders; this would then enable autonomous vehicles to follow conventional vehicles at short distances; clearly, it now needs to be debated if it is desirable to have such a vehicle at close distances in the rear mirror. Similarly, one needs to discuss the possible corresponding equipment (or not) of bicycles and pedestrians.[2]

Overall, it seems possible that driverless vehicles will need less flow capacity. Societal discussions are necessary to decide to which extent this should be used, and which equipment changes (if any) we want to accept for this.

7 Electrification

Simulation results for electrification are currently only available for the Berlin taxi fleet (Bischoff & Maciejewski, 2014, 2015). It should, however, be possible to transfer the results to a fleet of autonomous vehicles since the operating characteristics are similar. For the electrification of the taxi fleet we found that equipping about 300 of the overall existing Berlin taxi ranks with fast charging stations (50 kW) would be sufficient to fulfill the charging demand.

[2] Vgl. https://www.safewise.com/resources/wearable-gps-tracking-devices-for-kids-guide

At this point it is relevant to note that the fleet needs a certain amount of kWh per day, and they can be provided by charging stations of low or of high power. It turns out that prices for charging stations are roughly proportional to their power, and thus one has, at the same price, the choice between many slow or not so many fast charging stations. Here, the fast charging stations offer more flexibility for the same price; for example, when having multiple discharged vehicles, the fast chargers will send some of these vehicles back into operations after 1/2 h. For this reason, the investigation concentrated on fast charging stations. In cities where there is currently a sizable fleet of electric taxis operative, charging power is almost exclusively provided by fast charging infrastructure (Zou et al., 2016).

The simulations now need, besides the dispatch for the customers (Sec. 4), also a dispatch for the charging. Multiple approaches were considered; a simple one is that vehicles with a state of charging (SoC) of less than 30% when the customer gets off are assigned to a nearby charging station. It turns out that this is not a major problem: A taxi in Berlin typically travels 150 km per day, which can in principle even be done with a single charge. In Sec. 5, the vehicles travel 280 km per day in the average, but even here there is sufficient time between the peaks to recharge the vehicles (cf. Fig. 3). Therefore, it is sufficient to provide the necessary charging capacity spatially distributed in the system. Starting from these considerations we computed several scenarios (Bischoff & Maciejewski, 2015). The scenario that was most difficult for Berlin was one with increased demand (e.g., because of a trade show) in conjunction with low temperatures: At low temperatures, the charging capacity was reduced by a factor of two, while at the same time a factor of two more energy was necessary to heat the vehicles. Since the combination of these two effects concerns on average only a handful of days per year, our recommendation would be to rather install fossil heating, as is common, e.g., for caravans. In the longer run, one might also consider better vehicle isolation, or keeping the batteries warm. Furthermore, the use of heat pumps in vehicles has already increased the heating efficiency of modern electric vehicles. For our Berlin situation, air conditioning during warm days was less of a problem; this may be different for other geographical regions.

Such considerations eventually lead to Tab. 1, which shows that the annual operating costs of fully battery-electric fleet vehicles are similar to those of hybrid-electric vehicles using fossil fuels. This includes the energy costs of all electrical devices including heating and air conditioning, and for electric vehicles the cost of the charging stations infrastructure is divided by the number of kWh and added to the cost of each kWh. For the battery it was assumed that it would last for 100 000 kilometers. As result, one obtains that running a battery-electric fleet on electricity is not more expensive than running a hybrid-electric fleet

Table 1: Annual operating costs of a battery-electric vehicle (BEV) and a hybrid-electric vehicle (HEV). Source Bischoff & Maciejewski (2015)

	BEV	HEV
vehicle kilometers [km/year]	75 000	75 000
energy costs [Eu/year]	4 620	6 390
spare battery [Eu/year]	2 500	0
engine maintenance [Eu/year]	150	1 000
sum [Eu/year]	7 270	7 390

on fossil fuel. – There may be other computations that lead to other results. However, it is plausible to assume that both electricity prices and battery prices will fall, implying that the current approximate equality should move towards a cost advantage for the battery electric fleet.

Overall it is thus to be expected that the operational cost of a battery electric fleet is lower than that of a fossil fleet. Remaining are the investment costs for the transition. Since, however, the investment costs for the charging stations are already included in the calculation, the remaining items are the new vehicles. Since here the cost of the battery is already included into the operating costs, it remains the price of an electric car minus the battery. It is to be expected that this price will eventually be below that of a conventional vehicle, so that the conversion of the fleet could be funded through the regular replacement investments.

Thus, the emerging overall picture is that both the Berlin taxi fleet as well as a potential fleet for all of the inner-urban car traffic could be operated battery-electrically, with presumably lower costs than the current system. Admittedly, this statement assumes that the necessary raw materials, e.g. Lithium for the batteries, can be obtained in the necessary amounts at current prices; also, issues such as social or environmental standards during production are outside the scope of this argument. Both are outside our area of expertise.

8 Pooling

The text so far considered trips with maximally one passenger per vehicle. One can approach pooling via an algorithm as it is, e.g., used by Uber(Uber, 2017): Insertion of an additional passenger into an already assigned trip is accepted if for none of the passengers their trip becomes longer than

$$\alpha \times t_{single} + \beta \ , \tag{1}$$

where t_{single} is the trip time as single passenger. Typical values are $\alpha = 1.7$ and $\beta = 5\,min$; if, say, the single passenger trip time is $10\,min$, then the trip should not become larger though pooling than

$$1.7 \times 10 \ min + 5 \ min = 22 \ min \ . \tag{2}$$

The simulation described in Sec. 5 can still proceed as before, but the dispatch (Sec. 4) will now also consider to add further passengers to vehicles with passengers, as long as condition (1) is fulfilled for all passengers. This pooling algorithm is presented in more detail and applied to a data set with taxi requests of a weekday in Berlin in Bischoff et al. (2017). Figures 7 and 8 are taken from that study.

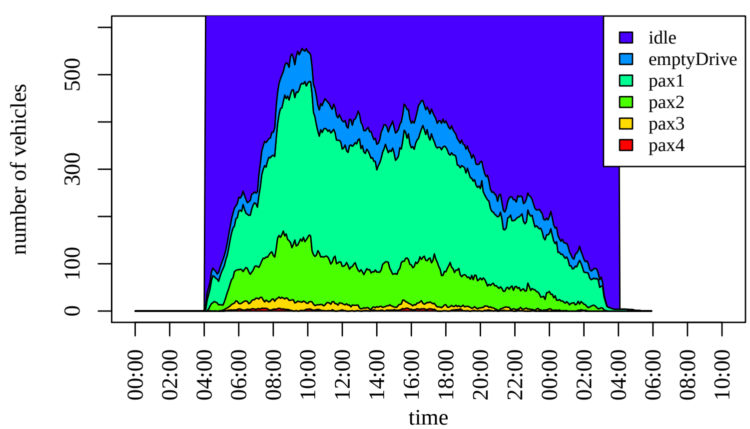

Figure 7: Vehicle occupancy as a function of time (Capacity 4 passengers, $\alpha = 1.7$ und $\beta = 2min$). Bischoff et al. (2017)

We typically obtain curves such as in Fig. 7. One observes that at all times about 40% of the occupied vehicles have more than one passenger. However, one also observes that these are usually just two passengers; three or more are rare even when we assume vehicles with no limit on seats. This corresponds to results by, e.g., Knapen (Knapen et al., 2013), which state that vehicle pooling is difficult to achieve if one does not want to accept long detours.

For the Berlin scenario, the average trip time increases from about 20 min to about 23 min (Fig. 8). The overall share of shared trips (with at least two

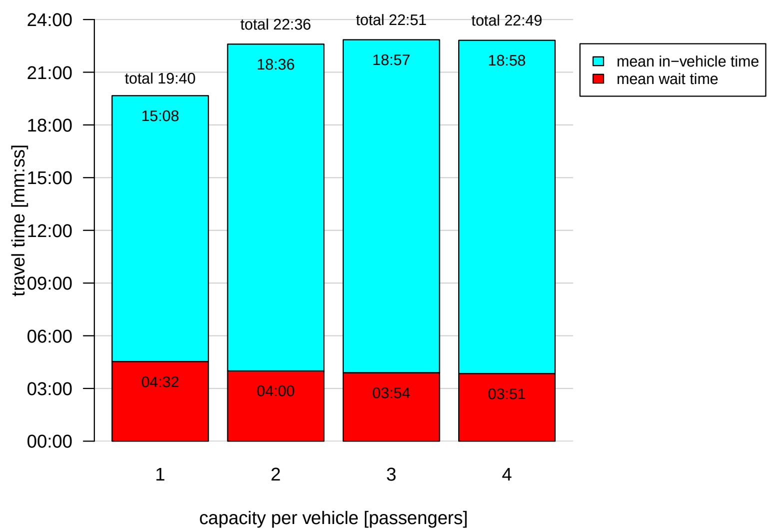

Figure 8: Warte- und Fahrzeiten der Sammeltaxis ($\alpha = 1.7$ und $\beta = 2min$). Vgl. Bischoff et al. (2017).

passengers for parts of a trip) is 61% in the scenario with four seats per vehicle. Pooling is more likely to happen on longer trips, as the probability of matching passengers increases. If only trips between two and five kilometers of direct travel distance are evaluated, then the total travel time of unshared trips is around ten minutes, whereas it is around 14 minutes for shared trips. According to equation 1, up to 22 minutes would be allowed. In how far such detours fit into the actual travel plans of each individual is up to their own decision. We suspect that many travelers may not find this attractive. At any rate, with this approach it is possible to fulfill the full Berlin innerurban car demand, and to end up with fewer vehicle kilometers than with the privately owned cars, so that such an approach would also be possible without the traffic flow gains through autonomous vehicles.

9 Shared autonomous vehicles as a replacement for conventional public transit

It is plausible to also consider fleets of shared autonomous vehicles as a replacement for conventional bus and tram lines. The following two studies allow some speculation.

9.1 Access to and egress from mass transit (the "last mile")

In our simulation one can replace the public transit (PT) router, which generates trips with access and egress by walk, by one which generates access and egress using pooled SAVs. The overall complexity of the simulation system increases, mostly because of the additional complexity of the simulation of schedule-based public transit, with the result that we do not have fully reliable simulations in this area yet.

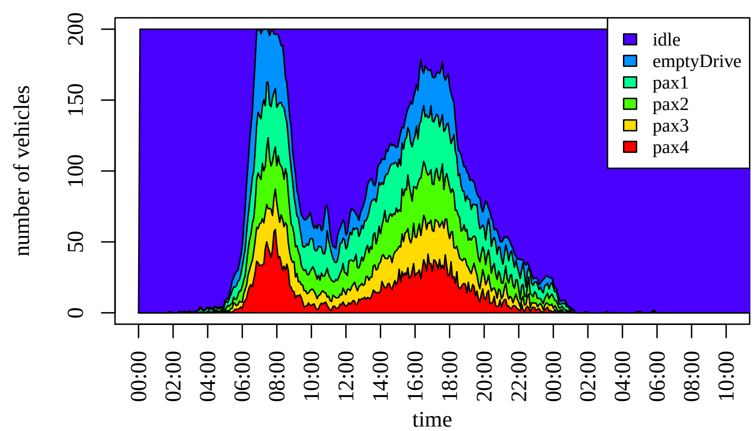

Figure 9: Vehicle occupancy for last mile study (section 9.1).

Gregor Leich (Leich, 2017) has done a study in which for a well demarcated study area in the northwest of Berlin (in the area of Heiligensee and Konradshöhe) the current bus supply was replaced by various constellations of SAVs. Connecting urban rail services were left unchanged. With respect to costs, the study compared driverless robotaxis with driverless conventional buses. For all investigated constellations, the conventional (but driverless) buses were slightly to significantly less costly. Evidently, the operator is able to reach a high loading factor for its conventional buses.

With respect to trip times, the picture is complex. In-vehicle times increase slightly, which can be explained by the additional trips of the robotaxis to and from individual addresses and by the detours to serve multiple ride requests with the same shared vehicle. Because of the door-to-door service, walking times decrease considerably. The overall trip times decrease by approximately 100 sec per trip. At the same time, the simulation at most times (except peak hours) has a large number of idle robotaxis. It should be investigated how the situation changes when these idle robotaxis are used to increase passenger benefits, possibly with the consequence of less pooling and therefore higher costs. Furthermore, fleet rebalancing could decrease robotaxi wait times, but was not available when the simulation was run.

9.2 Complete replacement in a medium-sized city

Bischoff et al. (2018) have conducted a simulation study for Cottbus, Germany, where all public transit (mostly buses and trams) was replaced by various constellations of SAVs. The study shows that operations without door-to-door service and instead serving designated stops leads to lower operating costs, but also to lower benefits. Both SAV approaches, without drivers, appear to have lower operational costs *and* lower passenger travel times than the current system, with drivers. The results are not directly comparable with the study by Leich since, for example, the current public transit vehicle loadings in Cottbus are rather different from those in Berlin, and public transit was completely removed, so there was no wait for connecting trains.

9.3 Implications on public transport

There are other studies with different backgrounds using other assumptions. For example, ISV (2016, pp. 42) have replaced all conventional cars and bus lines in the Stuttgart region by pooled SAVs with six seats. In one scenario, where rail-bound mass transit continues to operate, more than half of the departures between 6:00 and 19:00 are fully occupied. Pooling is thus assumed to work better. However, this can only partly be compared with the results presented in Fig. 9, as the assumptions differ: The Stuttgart study assumes a larger overall demand per area, taking both car and bus demand into account. Furthermore, passengers are pooled a priori, and the demand is bundled into traffic analysis zones and intervals of 15 minutes. In contrast, the MATSim models we use provide a more detailed resolution of passenger demand, and thus make it harder to pool passengers and the passenger demand only consists of former bus passengers. It will have to be investigated, to what extent spatial and temporal schedule deviations will be accepted by passengers.

So far, in our work we have not included effects of mode choice into their studies. Hörl et al. (2017) assume in a study for Zürich that pooled SAVs will be comparable in price with private cars, however they are likely to be more expansive than currently operative, subsidized public transport. As shown in the previous two sections, SAVs will not always lead to lower travel times compared to public transport. Fleet operators may vary their vehicle dispatch in multiple ways and experiment with different service levels. Nevertheless, also today customers of ride-hailing services often state that the travel time reduction is their main reason to use such services rather than public transport (Clewlow & Mishra, 2017).

Overall one learns, not fully unexpectedly, that an evaluation of SAVs as replacement of or addition to the conventional system of public transit is more complex than for the car system. One question is if the cost savings from possibly no longer having to pay for a driver should rather be invested into reduced prices/subsidies, or alternatively into better service quality. If the latter is desired, then particularly large improvements will be possible in areas where the conventional buses are not full even during the peak periods. Furthermore, a balance needs to be found between the cost savings by pooling on the one hand, and the travel time extensions that this causes. More studies, for different regions with different characteristics, need to be undertaken. Another topic to investigate are the effects of competition between the remaining conventional public transit system and the new fleet operators.

10 Discussion

10.1 Operator models, reservations, and prices

For the assessment of a robotaxi service the operating rules need to be specified. Our simulation currently assumes that travellers, when ending their activity, push some kind of button, which makes a robotaxi drive to the customer to pick her up. In the Berlin taxi system this is largely sufficient; waiting times of more than 5 min are rare, and normally do not justify the additional effort of a pre-booking. In thinly populated regions, however, this will not work, since available vehicles will be too far away without pre-booking. Pre-booking becomes indispensable when a certain arrival time, for example for a train station or an airport, is to be kept: Here an operator needs to reject requests that come in too late, leading to a latest pre-booking time. Rejection is also a possible approach to deal with an undersized fleet; alternatively, an operator could charge different prices for immediate pickup vs. some later pickup.

10.2 Preemptive vehicle rebalancing

Operators can leave vehicles at the locations where the last passenger left the vehicle, but they can also pre-emptively drive to areas where they expect high demand. This is, for example, plausible for commuter traffic. We sometimes hear the question if empty vehicles should "cruise" in the same way as taxis cruise. Our intuition is that this is only useful when there is also some "visual" channel for taxi requests, corresponding to the current flagging down of taxis. When, however, taxis are ordered via smartphone or similar, then the pre-emptive rebalancing should already have reached the optimum, and all additional vehicle movements will only generate costs, without reducing the expected trip time to the next customer.

10.3 Induced traffic and regulation

It is to be expected that such traffic systems will be attractive, and thus generate additional trips (= induced traffic). Similarly important will be mode choice reactions: persons who are currently walking, using the bicycle, or public transit, and who might switch to using an autonomous fleet. If one puts this together with a possible flow capacity increase, we might end up with twice as many vehicles on the road as today. It needs to be discussed if this is what we want, especially in urban situations. If not, an alternative might be a licensing scheme, where licenses are auctioned off to a small number of suppliers in regular intervals. Such an approach is, for example, used for the German regional trains. Such a license could include service requirements for the surroundings of the core city, similar to regulation in the area of telecommunication where the connection fee is always the same no matter where a property is located.

10.4 Possible transition scenarios

We concentrate on scenarios where we assume that the transition has already taken place, and has been complete. Clearly it is plausible to think about transition scenarios. Generally we would expect that autonomous vehicles may establish themselves from three different directions:

Scenario "Californian freeway" Freeways (with directional separation) are less complex than urban traffic, and safer than roads without directional separation. It is with certain vehicle types already now possible to take the hands off the steering wheel, even when currently the responsibility remains with the driver. Possibly problematic freeway sections could be improved through road markings or a small number of electronic navigation beacons.

Scenario "urban fleet" This corresponds to the scenario discussed most in the present paper: A fleet of autonomous vehicles which are either in addition to conventional vehicles, or replace them.

Scenario "rural fleet" A third possible scenario would be a robotaxi system that operates in thinly populated rural areas. Here, the autonomous vehicles would replace infrequent and under-utilized regular buses.

For each of these scenarios, a different vehicle type might be useful: for the "Californian freeway" maybe vehicles similar to the ones we have today, for the "urban fleet" probably vehicles that are smaller and more lightweight than today's, with smaller maximum speed, while for the "rural fleet" one might want minibuses that could also get larger groups home from the Friday night disco.

One might also imagine the following: the vehicle comes autonomously to the customer, the customer drives herself to the destination, the vehicle then continues autonomously to the next customer or to parking. Such an approach would have the advantage that the vehicle could drive slowly while in autonomous mode, and high speeds would only be reached with human drivers.

There is no consent when AV technology will be operative in a way that they could fully replace conventional cars. Some expect that vehicles will be available on the mass market from 2030 on, while others argue that AVs will initially not be able to cope with averse weather or road conditions. Thus, middle-class vehicles with full AV capabilities may only enter the market in the late 2030s or early 2040s. During the long transition phase many of the expected advantages of AVs in vehicle flow may not materialize or not materialize fully. Some first pooled SAV services are likely to be established in the 2020s and 2030s. Initially, their share on modal split might be rather low and many households will keep using private vehicles. This will allow operators to experiment with pricing policies and fleet management.

11 Conclusions

Overall, the situation is quite complex. For Berlin, we draw the following conclusions:

1. A switch of all private car traffic within the city limits to a fleet of driverless autonomous vehicles would in general be possible.
2. These vehicles will generate more vehicle kilometers than currently, since the empty trips from one dropoff to the next pickup come on top of the passenger trips. In our scenario, the increase is 13%.

3. Most researchers expect increased flow capacities for autonomous vehicles. The increase would have to be at least 13% to compensate for the additional empty trips alone.

4. The possible electrification of such a fleet is already at today's prices not more expensive than running it on fossil fuel.

5. Autonomous vehicles, presumably in the form of minibuses, are good candidates for access to and egress from mass transit, especially in areas where the conventional buses are running half-empty also during peak periods.

The following questions should be discussed by society:

- Do we want the to-be-expected significant increase of the number of vehicles on the road in the urban areas?
- What do we want to do with the parking spaces that will be freed up?
- Do we want to regulate fleet providers through licenses, for example to enforce a minimum service guarantee also outside the core city?
- Do we want that autonomous vehicles offer arbitrary point-to-point connections, or should they concentrate on access to/egress from mass transit?

References

Ahrens, G.-A., F. Ließke, R. Wittwer, S. Hubrich, & S. Wittig. Mobilität in Städten – SrV 2013. 2014. URL http://www.tu-dresden.de/srv/.

Axhausen, K. W. Sind wir bald alle Passagiere? Neue Züricher Zeitung, 4.10.2016, 2016.

Bischoff, J., K. Führer, & M. Maciejewski. Impact assessment of autonomous DRT systems. VSP Working Paper 18-04, mobil.TUM 2018, 2018. URL http://www.vsp.tu-berlin.de/publications.

Bischoff, J. & M. Maciejewski. Agent-based simulation of electric taxicab fleets. *Transportation Research Procedia*, 4, 191, 2014. doi:10.1016/j.trpro.2014.11.015.

Bischoff, J. & M. Maciejewski. Electric taxis in Berlin – Analysis of the feasibility of a large-scale transition. In J. Mikulski, ed., *Tools of Transport Telematics*. Springer International Publishing, 2015. doi:10.1007/978-3-319-24577-5.

Bischoff, J. & M. Maciejewski. Autonomous taxicabs in Berlin – a spatiotemporal analysis of service performance. *Transportation Research Procedia*, 19, 176, 2016a. doi:10.1016/j.trpro.2016.12.078.

Bischoff, J. & M. Maciejewski. Simulation of City-wide Replacement of Private Cars with Autonomous Taxis in Berlin. *Procedia Computer Science*, 83, 237, 2016b. doi:10.1016/j.procs.2016.04.121.

Bischoff, J., M. Maciejewski, & K. Nagel. City-wide shared taxis: A simulation study in Berlin. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2017. doi:10.1109/itsc.2017.8317926.

Clewlow, R. & G. S. Mishra. Disruptive transportation: The adoption, utilization, and impacts of ride-hailing in the United States. Tech. rep., Institute of Transportation Studies, University of California, Davis, 2017. URL https://itspubs.ucdavis.edu/wp-content/themes/ucdavis/pubs/download_pdf.php?id=2752. Last accessed 12.04.2018.

Flötteröd, G. & B. Kickhöfer. Choice models in MATSim. In Horni et al. (2016), chap. 49. doi:10.5334/baw.

Horni, A., K. Nagel, & K. W. Axhausen, eds. *The Multi-Agent Transport Simulation MATSim*. Ubiquity, London, 2016. doi:10.5334/baw.

Hörl, S., C. Ruch, F. Becker, E. Frazzoli, & K. W. Axhausen. Fleet control algorithms for automated mobility: A simulation assessment for Zurich. Annual Meeting Preprint 18-02171, Transportation Research Board, Washington, D.C., 2017.

Infas & DLR. Mobilität in Deutschland 2008 Ergebnisbericht. 2010.

ISV. Modellergebnisse geteilter autonomer fahrzeugflotten des oeffentlichen nahverkehrs. Tech. rep., Institut für Straßen- und Verkehrswesen, Universität Stuttgart, 2016. URL https://www.isv.uni-stuttgart.de/vuv/publikationen/downloads/MEGAFON_Abschlussbericht_V028_20161212.pdf.

Knapen, L., D. Keren, A.-U.-H. Yasar, S. Cho, T. Bellemans, D. Janssens, & G. Wets. Estimating scalability issues while finding an optimal assignment for carpooling. *Procedia Computer Science*, 19, 372, 2013. doi:10.1016/j.procs.2013.06.051.

Kuhn, H. W. The Hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2), 83, 1955.

Leich, G. *Agentenbasierte Simulation autonom fahrender Rufbusse im Nordwesten Berlins*. Master's thesis, TU Berlin, Institute for Land and Sea Transport Systems, Berlin, Germany, 2017.

Levinson, D. M. & K. J. Krizek. *The end of traffic and the future of transport*. Amazon Digital Services LLC, 2nd ed., 2015.

Maciejewski, M. Benchmarking minimum passenger waiting time in online taxi dispatching with exact offline optimization methods. *Archives of Transport*, 30(2), 67, 2014a. doi:10.5604/08669546.1146978.

Maciejewski, M. Online taxi dispatching via exact offline optimization. In *7th International Conference on Logistic Systems – Theory and Practice*. 2014b. Also VSP WP 14-14, see http://www.vsp.tu-berlin.de/publications.

Maciejewski, M. & K. Nagel. The influence of multi-agent cooperation on the efficiency of taxi dispatching. In R. Wyrzukowski et al., ed., *10th International Conference Parallel Processing and Applied Mathematics (PPAM) Part II*, no. 8385 in LNCS. Warsaw, Poland, 2013a. doi:10.1007/978-3-642-55195-6_71.

Maciejewski, M. & K. Nagel. A microscopic simulation approach for optimization of taxi services. In T. Albrecht, B. Jaekel, & M. Lehnert, eds., *3rd International Conference on Models and Technologies for Intelligent Transportation Systems 2013*, pp. 1–10. TUDpress, 2013b. Also VSP WP 13-12, see http://www.vsp.tu-berlin.de/publications.

Uber. uberPOOL – Sharing is saving, 2017. URL https://www.uber.com/nyc-riders/products/uberpool/. Accessed 12.05.2017.

Zou, Y., S. Wei, F. Sun, X. Hu, & Y. Shiao. Large-scale deployment of electric taxis in Beijing: A real-world analysis. *Energy*, 100, 25, 2016. doi:10.1016/j.energy.2016.01.062.