Autonomous vehicle fleet sizes required to serve different levels of demand

Required Autonomous Vehicle Fleet Sizes to Serve Different Levels of Demand

Patrick M. Boesch, Corresponding Author
IVT, ETH Zurich
CH-8093 Zurich, Switzerland
Tel: +41 44 633 39 52; Fax: +41 44 633 10 57; Email: patrick.boesch@ivt.baug.ethz.ch

Francesco Ciari
IVT, ETH Zurich
CH-8093 Zurich, Switzerland
Tel: +41 44 633 71 65; Fax: +41 44 633 10 57; Email: ciari@ivt.baug.ethz.ch

Kay W. Axhausen
IVT, ETH Zurich
CH-8093 Zurich, Switzerland
Tel: +41 44 633 31 05; Fax: +41 44 633 10 57; Email: axhausen@ivt.baug.ethz.ch

Word count: 6,121 words text + 5 tables/figures x 250 words (each) = 7,371 words

TRR Paper number: 16-2682

Submission Date: March 03, 2016

ABSTRACT
Automated Vehicles (AV) promise many benefits for future mobility. One of them is a reduction of the required total vehicle fleet size, especially if AVs are used predominantly as shared vehicles. This paper presents research on this potential reduction for the greater Zurich region, Switzerland. Fleets of shared AVs, serving a predefined demand, are simulated with a simulation framework introduced in the paper. Different scenarios are created, combining different levels of demand for AVs with different levels of supply (i.e. AV fleet size).

An important contribution of this study is the use of a spatially and temporally highly detailed travel demand, going beyond the simplifications of previous studies on the topic. This provides a more solid basis to the ongoing discussion on the future fleet size.

It is found that, for a given fleet performance target (here 95% of all transport requests are served within 5 minutes), the relationship between served demand and required fleet size is non-linear and the ratio increases as demand increases. There is a scale effect, which has the important implication that for different levels of demand the fleet is used more or less efficiently.

This paper also finds that, if waiting times of up to 10 minutes are accepted, a reduction of up to 90% of the total vehicle fleet can be possible even without active fleet management like vehicle redistribution. Such effects require, however, that a large enough share of the car demand can be served by AVs.

1   INTRODUCTION

The idea of an Automated Vehicle (AV) is already several decades old. For a long time, research on the topic addressed exclusively the technological aspects of AVs. Recent developments, however, made clear that the technology will soon be available. Car manufacturers and new players are already testing and selling vehicles which, to various degrees, can be considered autonomous. To clarify the difference among automation levels a scale has been proposed by NHTSA (*1*). At the top of the scale, level 4, are vehicles that can drive without a human driver in uncertain, open environments created for humans. Such vehicles are expected to arrive in the consumer market within the next few years (*2*). The consequences on the transportation system are yet uncertain, but it is reasonable to assume that they potentially will be extremely far reaching. One fundamental question is, how AV usage will be regulated. At the moment AVs require special test permits to be operated in open traffic conditions. It is clear that this restriction will be lifted at some point, but the form in which AVs will be allowed is obviously a key point. Some researchers pointed already to the fact that to effectively harvest all the possible benefits of AVs, a "car-sharing" like scheme should be preferred to an ownership based model. AVs would overcome most of the limitations that car-sharing services have versus personal ownership and usage. For example, compared to car-sharing in any of the currently available forms, with AVs one would not need to travel to the next vehicle. Instead the vehicle would be able to drive exactly where the demand is, much like a taxi. The hypothesis is that a car-sharing scheme would provide the same kind of flexibility that cars provide, without (or limiting) the possible negative impacts of the widespread diffusion of AVs. Specifically, thinking of a large scale AV car-sharing scheme able to satisfy most of the demand currently covered by personal cars, the most evident benefit would be the much smaller total car fleet needed. It would enable a more efficient use of the vehicles in terms of productive time with a cascade of positive outcomes, for instance reduced parking requirements and thus freeing up large areas of high value urban space.

A few papers have already been published hypothesizing AV car fleets offered in shared vehicle schemes. These studies find that current travel demand could be served with fleet size reductions of up to 90% of the current car fleet sizes (*3*).
For example (*4*) and (*5*) investigated the required fleet size to serve the current demand in Singapore with an "Automated Mobility-on-Demand System". Using queue models, they found that one third to 40% of the current number of passenger vehicles in Singapore would be sufficient to serve the total travel demand of Singapore.
Other researchers approached the fleet size problem with simulations. In general they are characterized by simplifications of the demand in terms of temporal resolution, spatial resolution or both compared to current state-of-the-art in simulations of transport.
In (*3*) and (*6*), on the temporal axis, travel demand is randomly generated using transport demand and trip duration patterns. The spatial distribution is rather coarse (grid based in the case of (*3*) with a grid cell size of 0.25x0.25miles) and random. In (*3*) the travel times to calculate the re-availability of vehicles are based on Manhattan distances and average speeds. They find that 10% (*3*) or 15% (*6*) of the current fleet would be enough to serve the current demand.
In (*7*) and (*8*) the demand is explicitly generated with demand models to correctly represent the real demand of the respective area (New Jersey (*7*) and Lisbon, Portugal (*8*)). The subsequent transport simulations work with static demands, which means that the demand can not adapt to the available transport supply and traffic situations. Travel demand is aggregated to zones after generation (zone size 0.5x0.5 miles (*7*) and 200x200 meters (*8*)). While (*7*) does not consider travel times and thus re-availability of vehicles (they provide an infinitely sized fleet), in (*8*) travel

times are calculated by routing the trips on a network and assuming average link occupancies and thus speeds by time of day. In (*7*), they find that a substantial ride sharing potential exists for AV fleets. For Lisbon (*8*) suggests a possible fleet size reduction of 90%.

A recent study by Chen et al. (*9*) investigated how a fleet of "Shared, Electric, Autonomous Vehicles" performs if vehicle charging is considered too. They find that one vehicle could replace 3.7 to 6.2 privately owned vehicles.

This paper presents an investigation of the fleet size problem based on a simulation approach too, but with some important novelties. A fleet of shared AVs, which serves a predefined demand, is simulated. Several simulations are run, assuming that different shares of the current car travel demand are fulfilled with differently sized AV fleets.

The demand represents the car travel demand for the greater Zurich region, Switzerland with high spatial and temporal resolution. The demand and the travel times considered in the AV simulation were obtained using MATSim, an activity-based multi-agent transport simulation framework (*10*). This MATSim generated demand is used as a static demand in the AV simulation.

This paper introduces a simulation framework which goes beyond those used so far in the existing literature in terms of spatial and temporal resolution of the demand. This provides a more solid basis to the ongoing discussion on the fleet size required to serve a certain travel demand with a given level of service. It is argued here that precise representations of the start time and location, as well as of the destination and the expected travel time of a transport request in realistic traffic situations, are important aspects which need to be considered. Real transport demand shows complex spatial and temporal patterns and is tightly dependent on transport supply. If transport supply limitations are reached and congestion builds, travel times will increase nonlinearly and wider network effects will occur. The demand generation with the agent-based transport model MATSim and its equilibrium approach ensures that these effects are realistically represented. As a first application, the simulation framework was used to provide, compared to existing literature, a more systematic answer to the question of the required AV fleet size to serve different levels of demand. In total 140 different scenarios were simulated to thoroughly investigate the question. It is found that, for a fixed level of service, the relationship between trips served and fleet size is non-linear and the ratio increases as the number of trips increases. So there is, as could be expected, a scale effect, which has the important implication that to different levels of demand the fleet is used more or less efficiently.

This first application also allows to compare the scenario results with existing studies on the topic. Compared to earlier studies (i.e. (*3*) and (*8*)) a much lower peak fleet usage was achieved. The average usage is lower too as the trips served per day per AV and the total usage time of an AV is lower. While some of these effects can be explained by the larger scenario area and the lower population shares used here, at least some of this difference is also attributed to the more detailed demand. This is supported by the fact that still only 10% of the existing car fleet is sufficient to serve the current car travel demand, which also hints to a strong influence of the spatio-temporal characteristics of the demand on the possible substitution rate.

2   THE SIMULATION FRAMEWORK

On top of MATSim a separate simulation framework was developed. It simulates an AV fleet that serves trip requests. Travel demand and travel times are generated beforehand with a MATSim simulation.

2.1   Generating travel demand with MATSim

To understand the simulation framework, a basic understanding of MATSim is required. MATSim is an activity-based, multi-agent transport simulation framework (*10*). In MATSim agents optimize their mobility behavior in a co-evolutionary process until a dynamic user equilibrium (DUE) is reached. The agents represent a real population and have daily plans listing activities they would like to perform at certain times and locations. Transport demand arises from the execution of these plans within a realistic representations of transport supply. Single vehicles' movement is modeled through a queue based simulation on a highly detailed network. The temporal resolution of the simulation is one second. The initial population of agents, i.e. the initial set of daily plans, is created from travel diary survey data and population statistics (*11*). In the co-evolutionary process the agents can try different modifications of their plans (e.g. rerouting, rescheduling etc.) over many iterations until a DUE is reached. A scenario is calibrated such that traffic counts and transport statistics match real data. In the end, this results in a plausible travel demand representing an average working day in the simulated area.

2.2   Overview of the Simulation Framework
In the AV simulation, requests are made based on the MATSim-created demand. A request is assigned to the closest, available AV within a predefined service radius. This assigned AV moves to the start location and – if arriving early – waits for the passenger to finish its activity. Once the AV and the passenger are both ready, they move together to the trip end location. After they arrive, the AV becomes available for a next request. Movements are modelled as direct relocations with delayed arrival since a network representation is not part of the current implementation. For the movements to the request locations, travel times are calculated from corrected beeline distances and average speeds. For the actual trips travel times from MATSim are used. The MATSim travel times represent the expected travel times in realistic traffic situations. In the following sections the simulation framework is presented in more detail.

2.3   Initial Distribution of AVs
The initial distribution of AVs can have a strong influence on the simulation results and is therefore an important step. Different distribution approaches are suggested in the literature. For example the creation of new AVs at the locations of the first *n* requests in the simulation, with *n* being the target number of AVs ("warm-start" as in (*3*)), or a random distribution of the AVs in the study area (*6*).
To initially distribute the AVs, here the simulation framework first randomly samples for each AV one agent from the demand and then places the AV at the home location of this agent.This approach assumes that the sampled agent was the one who last used the shared AV the day before. This procedure determines only the initial locations of the AVs and all AVs are available to all agents from the start.

2.4   A Typical Trip
This section explains the characteristics of the simulation framework by following one request step by step from start to end.

2.4.1   Registering of a New Request
A request represents a trip. A trip consists, among other information, of a start location, a start time, an end location, and of the arrival time found in MATSim. Five minutes before the start time of a trip, that is five minutes before the previous activity ends, the request is registered. This means

assuming that people expect an AV to arrive within five minutes after they register a request. A request is served "*in time*" if the AV arrives within these first five minutes. If the AV arrives in the first five minutes after the activity ended, that is five to ten minutes after the request was registered, the request is served "*late*". This choice of 5 minutes for "*in time*" and 10 minutes for "*late*" is based on a study by Fagnant and Kockelman (*3*). If no AV can serve the request within 10 minutes, the request is "*unserved*" and dropped.

2.4.2   Search for a Free AV and Assignment

After a request is registered, an available AV is searched. The search process consists of first identifying the closest AV to the request location and then checking if this AV is within the service radius of the request. The service radius is the driving time radius equal to the remaining time until time out (10 minutes after registering (section 2.4.1)). If no suitable AV can be found, the search is repeated every simulation step until time out. If a suitable AV is found, it is assigned the request and the assignment time is recorded. An assigned request will not be reevaluated for closer AVs.

2.4.3   Drive of the AV to the Start Location

An assigned AV moves directly to the start location of the request. The duration of this movement, and thus the arrival time of the AV at the start location, is determined by first multiplying the beeline distance to the target location with a beeline factor and then dividing the product by the average speed of a car. The beeline factor, 1.43, was computed by comparing for all car trips in the MATSim demand the street network distance with the beeline distance. The average speed, 11.28 m/s, was calculated from the 2010 travel diary data for the greater Zurich area (*12*). If the AV arrives at the location of the request before the start time of the request (end time of the preceding activity of the agent), it has to wait until the start time.

2.4.4   Service

Once the AV has arrived the start location and the start time is reached, the request is served and the AV moves to the end location of the trip. The trip duration from MATSim is used to calculate the arrival time. The simulation therefore uses realistic travel times under realistic traffic conditions on the network at the time of the trip. Upon arrival the AV is released and ready for a next request.

2.5   Other Specifications

The simulation is a time-based simulation. This means that simulation steps are executed every fixed time interval (1 second for the study presented in this paper).
Following the standards of the MATSim framework, the simulation simulates one typical week day from midnight to 6am the following morning (30 hours). Statistics are sampled every 5 minutes of simulated time (statistical sample interval).

3   SIMULATED SCENARIOS

The case study is a MATSim scenario of the region of Zurich. It is a well-tested scenario used in previous studies (e.g. (*13*) or (*14*)). It simulates around 2.1M agents. They represent all persons that are either living and/or carrying out activities in or passing through the area. This practically means that trip requests can come from anywhere in Switzerland and neighboring areas of bordering countries (Figure 1).

Different levels of AV travel demand were generated drawing samples from the set of all car users in the scenario, which amounted to 1.3M agents with a total of 3.6M trips. The levels ranged between 1% and 10% (in 1% steps) of the total car travel demand. These ten demand levels were combined with ten different supply levels. For each supply level, the number of provided AVs equaled a different percentage (10% to 100% in 10% intervals) of the sampled users. This resulted in 100 different scenarios.

As the sampling of agents for the scenarios and the initial distribution of the AVs in the simulation (section 2.3) are random processes, the above scenarios were all repeated ten times with different random seeds (which command the underlying algorithm). This resulted in 1'000 initial simulations.

After a first analysis of the initial simulations, additional supply levels were considered to refine the results in critical value ranges. The additional supply levels were 5%, 15%, 25% and 35%. These were again combined with all demand levels and repeated with all random seeds, resulting in an additional 400 simulations.

For the current study, two main assumptions were thus made. The first is that in an early phase only a fairly small percentage of the population would use AVs. The second is that only trips made with cars would be substituted by AV trips (no modal shift). Please note that these are extremely strong assumptions in terms of behavioral response to the introduction of AVs, but the representation of such response is beyond the scope of this exercise (discussion in section 5.1) which rather focuses on AV fleet's size given a certain demand.

4 RESULTS

4.1 Performance of the AV Fleets

Figure 2A shows the performance of the AV fleets responding to requests within 5 minutes. Each scenario represents the average across all random seeds. The standard deviations are not visualized since they are very small (average 0.15%, maximum 0.92%). This suggests a robustness of the simulation results against changes in the participating population.

The performance of the AV fleet is almost independent of the population size for the demand levels 4% to 10%. For the demand levels 1% to 3%, however, the performance drops. AV fleets of the size of 10% to 30% of the original car fleet perform in these scenarios considerably less well than in the other scenarios. For the same performance larger fleets are required.

4.2 Fleet Required to Serve 95% of the Requests in Time

In Figure 2B the lines represent a crosscut through Figure 2A. The darker line shows for different levels of demand the minimally required fleet size to serve at least 95% of the requests in time. The lighter line represents the same values if requests served late are included too. The values shown in Figure 2B were calculated with a linearization between the two simulated supply levels closest to 95%.

Figure 2B confirms the conclusions drawn from Figure 2A. For small demand levels the required minimal AV fleet to achieve a 95% performance is strongly dependent on the share of population participating in the program. The ratio between the original fleet and the AV fleet grows as demand density grows, meaning that the fleet is used more efficiently. For larger demand levels – 5% to 6% and more – the required ratio stabilizes. If the requests served late are considered too, smaller AV fleets are required for the same performance.

4.3   The 10% Demand – 10% Fleet Scenario
For the remaining part of the Results section, one scenario is presented in more detail. The 10% demand / 10% fleet scenario was chosen as an example because it combines a large demand with a small fleet which introduces competition for AVs in the scenario.

4.3.1   Overview
The scenario consists of 130'831 car-driving agents, which undertook a total of 366'124 car trips. These trips represent the demand for the AV fleet. In total 13'083 AVs were placed in the scenario, which is 10% of the number of agents.
Over all random seeds, these AVs served on average 257'000 requests (70.3% of the total demand) in time (in under 5 minutes) and 95'000 requests (25.9%) late (within 5 to 10 minutes). This left 14'000 requests (3.8%) unserved. Served requests (in time and late) were assigned a vehicle on average after 26 seconds. The average response time, that is the time for an AV to appear at the start location of the request (including assignment), was 3.11 minutes (driving time: 2.68 minutes). For trips served in time, the AV had to wait on average 3.56 minutes for the passenger to appear. For trips served late, the passengers had to wait on average 2.67 minutes for the AV to appear. The average trip duration of a served AV trip was 18.66 minutes. In comparison, the average trip duration of a car trip in the Zurich region is 19.70 minutes (*12*).

4.3.2   Throughout the Day
The average numbers above describe the scenario, but to understand the system behavior, the dynamics over time need to be investigated. Figure 3A presents the available vehicles and the vehicles in use over the course of a simulation day. It also presents the handled requests. These are all served or dropped requests within the preceding sample interval (5 minutes). The requests are presented in more detail in Figure 3B. It shows over the course of a simulation day how many requests could be served quickly (in less than 1 minute), in time, late or had to be dropped. Vehicles are either in use or available. Since no other state is allowed in the simulation, they always sum to the total fleet size (Figure 3A). The vehicles in use suffice to match the demand over the day, except for the morning peak (Figure 3A). This means that the morning peak cannot be fully served, which is confirmed by a rise of unserved requests at this time (Figure 3B). In contrast, the higher afternoon peak can be served by the vehicles.
For most time of the day, a large share (a third and more) of the vehicles are not in use (Figure 3A). Even at the time of the highest peak, 4'033 vehicles (30.8% of the total AV fleet) remain unused.

4.3.3   Requested Trip Distances
Figure 4 presents a histogram of the requested trip distances. For trips of up to 40 kilometers a service rate of more than 95% was achieved (Figure 4B). These trips are 90.5% of the total number of trips (Figure 4A). For trips with longer distances the service rate drops with increasing distances. For the 1.8% of trips which are longer than 150km the service rate is 68%.

4.3.4   Usage of AVs
Over a day each vehicle serves in average 26.9 (standard deviation 9.8) requests (in time or late). Figure 5 presents a histogram of the number of served requests per AV.
To serve these requests, an AV spends 4.0% of the day to pick-up agents and 27.9% serving. These two sum to a total productive AV time of 31.9% of the day. For the rest of the day, the vehicles are

either waiting for agents (3.9%) or waiting for a next assignment (64.2%). In comparison, in Switzerland today cars are used productively for only 3.2% of the day (calculated from (*15*)).

5   DISCUSSION

5.1   Simulation Framework
The presented framework is capable of simulating a fleet of AVs serving a spatially (meter-fine coordinates) and temporally (1 second intervals) detailed demand. This demand is generated with MATSim, an agent-based transport simulation framework, and used statically in the AV simulation. MATSim creates the demand with an equilibrium approach considering the transport infrastructure as well as system interactions. This process makes the travel demand, and in particular travel times, much more realistic than the randomly generated demand and the average speeds used in previous works on the topic ((*3*), (*4*), (*5*), (*6*)). Zachariah et al. (7) use a sophisticated demand generation too, but lack realistic travel times. The study on Lisbon (8) uses a detailed demand generation, but aggregates the demand to zones. Additionally, it doesn't use a dynamic traffic model like MATSim to calculate travel times, but use hourly updated network loads.
From a performance standpoint, the AV simulation presented here has short computation times, thanks to pre-calculation of demand and travel times (few minutes for small scenarios and up to few hours for the largest scenarios single-threaded on a standard desktop computer). This makes possible the investigation of many different combinations of supply and demand levels within hours.

There are, however, also some limitations to the framework. The most important limitation is probably that no redistribution was simulated. The ability to re-equilibrate system unbalances autonomously is one of the great advantages which are expected from AV car-sharing services over the existing ones. Attempts to include redistribution in the simulation revealed that a successful redistribution algorithm becomes complex to implement and calibrate for spatio-temporally detailed simulations. The redistribution algorithm not only has to determine where precisely idle AVs should be sent, but also how often a redistribution is meaningful for what share of the idle AVs and how to select the AVs which will actually be redistributed. Therefore, for this paper, it was decided to focus on the scenarios without redistribution. This is still an important step since it represents a baseline for future investigations of redistribution algorithms.
A second limitation of the simulation framework is that no network routing is currently implemented. Since no redistribution was simulated, however, only the short pick-up trips of AVs to the trip start locations were affected by this limitation. Inaccuracies should therefore be limited. Nevertheless, a network routing implementation is planned for the longer redistribution trips.
A third limitation is the static demand. With static demand, agents have no possibility to react to the new offer of an AV sharing service by changing their behaviour and mobility demand. Traveling with an AV might be very comfortable and thus induce additional demand and mode changes from non-car modes to AVs. This additional traffic, however, might increase travel times. The trade-off between these two effects and how it influences demand can only be investigated with more sophisticated transportation simulation frameworks allowing for a dynamic representation of travel demand (e.g. an integration of this framework into the MATSim iteration process or a direct simulation of AVs in MATSim).

5.2   Simulation Results

Compared to the literature, the 10%/10%-scenario (section 4.3) shows a much lower fleet usage. While here a peak usage of around 70% was achieved, it was 97% in (*3*) and 99.6% in (*8*). The same applies for the average usage. Here in the 10%/10%-scenario AVs are used about one third of the day (section 4.3.4) and serve in average 26.9 trips. In (*8*) they are used two thirds of the day and in (*3*) they serve an average of 35.87 trips per day.

It is argued that, besides relocation, these differences are mainly due to the inner city character of the other two studies. Inner cities show an above average high density of demand which results in very high usage numbers of the fleets. The demand used for this study, originating from locations over a very large area (more than 420km x 270km as shown in Figure 1), shows what should be considered if a more open scheme would be applied. Indeed, Figure 4, which shows the number of met requests versus trip distances, indicates that a limitation of the service area has strong influence on the average fleet usage and the level of service achievable.

It is interesting, however, that even for this large area, 1 AV has the potential to replace up to 10 of today's cars if a maximum reaction time of 10 minutes is accepted and the share of the participating population is large enough (Figure 2). Even though the ratio drops to 1 AV for 4 cars if only requests served within 5 minutes are considered, this result is very surprising because it was achieved without any redistribution. Given the discussion above, it is hypothesized that the high density of requests close to Zurich overcompensates the effects of the large area. Apparently a strong spatial aspect of the trip distribution comes into play. This is supported by the fact that most unserved requests happen during the morning peak, while the higher afternoon peak could be fully served (Figure 3). This means that, even if the results of this study, at a first glance, do not depart substantially from previous studies, it is important to be able to capture such effects which depend on a realistic and detailed demand. This relationship between city form (and density) and the efficiency of an AV fleet is a topic for future research. For the Zurich scenario, however, Figure 2 suggests that a limit is reached and that redistribution or carpooling will be required to achieve further improvements in fleet usage. Santi et al. (*16*) show for example that for taxi fleets large improvements can be achieved with optimized carpooling. Balac and Ciari (*17*) showed that improvements through redistribution can – at least for one way station based car-sharing – be substantial. How and how much the situation can be improved with carpooling and redistribution for AV car-sharing fleets, as well as what algorithms to use – (*18*) provides a recent review on car-sharing redistribution approaches – will be the topic of future research.

An important finding, despite the discussion above, is that the minimal AV fleet required to serve a demand successfully is dependent on the size of the demand (sections 4.1 and 4.2). For example, if 10% of the population participate, a service level of 95% can be achieved with 1 AV per 10 people. If only 1% of the population participates, 1.5 vehicles are required per 10 people (Figure 2B), which is 50% more. This can be explained by the available number of vehicles per area. The same fleet size (in relative terms) results in absolute numbers in less vehicles per area for smaller participating populations. If fewer vehicles are available per area, the chance for a request to find an AV within the service radius in time is smaller and thus more vehicles are required. It is interesting that this effect saturates if 5% or more of the total demand are simulated (Figure 2). What effect redistribution has on this relationship of fleet size on demand size is not clear yet. It is hypothesized that redistribution will be able to reduce this effect – to what extent and cost is part of future investigations.

6   SUMMARY AND NEXT STEPS

This paper presented a framework for the simulation of Automated Vehicles (AV). The framework builds on MATSim, a state-of-the-art multi-agent transport simulation, although it is not directly an extension of it. The framework uses the demand generated by MATSim, but simulates the AVs in a separate simulation environment. This demand is – to the best knowledge of the authors – spatially and temporally more detailed than the demand used in any other existing study on the topic, with some important implications as discussed in section 5. Despite the level of detail, computation time is short enough that series of large scale simulations (10^6 agents) are possible within hours. This performance allowed a detailed (140 simulated scenarios) investigation of how the demand size influences the required fleet size. A decreasing trend for the ratio vehicles/users was found as the number of users increases. This means that for a small demand the fleet is used less efficiently. If this observation can be confirmed by future research, it has implications not only for the transferability of results from studies of AV fleets but also for future experiments on shared services. If only small samples of the population are included, operators need to be aware that they will have to provide comparatively more vehicles. It was shown that the ratio of vehicles per participants increases more than 50% if the population share drops.

The simulation framework still has several limitations and future work will address them. The implementation of a redistribution scheme and of network routing are two of the most important next steps (section 5). For the research on AV fleets, however, it is probably more important to overcome first the limitation of using a static demand. Transport demand is influenced by new mobility offers and changed mobility costs. Demand has to be able to adapt to these new offers. AVs might induce more demand because they make traveling more comfortable and less expensive. This might also induce mode changes from public transport and slow modes to AVs (*19*). Such demand changes might bring the transport infrastructure to its limits and increase travel times with AVs. This in turn would have demand reducing effects, bringing the system to a new equilibrium. How these effects influence each other, how the new equilibrium looks like and what this means for infrastructure capacity requirements, requires the investigation with simulation frameworks able to explicitly account for such effects.

MATSim is inherently capable to model such complex interactions, assuming that the simulation of AVs is directly embedded in the framework too. For now MATSim was only used to generate the initial demand since it is not yet able to simulate AVs. After this is implemented (*20*), however, it will be possible to get a more complete picture of the impact of AVs on the global vehicle fleet.

7   ACKNOWLEDGEMENTS

REFERENCES

1.   National Highway Traffic Safety Administration. *Preliminary Statement of Policy Concerning Automated Vehicles*. Released May 30, 2013. http://www.nhtsa.gov/About+NHTSA/Press+Releases/U.S.+Department+of+Trans portation+Releases+Policy+on+Automated+Vehicle+Development. Accessed Jul. 28, 2015.

2.   Lutin, J. M., A. L. Kornhauser and E. Lerner-Lam. The revolutionary development of self-driving vehicles and implications for the transportation engineering profession. *ITE Journal*, Vol. 83, No. 7, 2013, pp. 28-32.

3.   Fagnant, D. J. and K. M. Kockelman. The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios. *Transportation Research Part C: Emerging Technologies*, Vol. 40, 2014, pp. 1–13.

4.   Spieser, K., K. Treleaven, R. Zhang, E. Frazzoli, D. Morton and M. Pavone. Toward a systematic approach to the design and evaluation of automated mobility-on-demand systems: A case study in Singapore. In G. Meyer and S. Beiker (eds.) *Road Vehicle Automation*, Lecture Notes in Mobility, Springer International Publishing, Berlin, 2014, pp. 229-245.

5.   Zhang, R., K. Spieser, E. Frazzoli, and M. Pavone. Models, Algorithms, and Evaluation for Autonomous Mobility-On-Demand Systems. Presented at *American Control Conference*, Chicago, 2015.

6.   Burns, L. D., W. Jordan and B. Scarborough. Transforming personal mobility. *Technical Report*, The Earth Institute, Columbia University, New York, 2013.

7.   Zachariah, J., J. Gao, A. Kornhauser and T. Mufti. Uncongested mobility for all: A proposal for an area-wide autonomous taxi system in New Jersey. Presented at the *93rd Annual Meeting of the Transportation Research Board*, Washington, 2014.

8.   International Transport Forum. *Urban Mobility System Upgrade – How shared self-driving cars could change city traffic*. Organisation for Economic Co-operation and Development, 2015.

9.   Chen, T.D., K. Kockelman and J. Hanna. Implications of a Shared, Autonomous, Electric Vehicle (SAEV) Fleet. Presented at the *14th International Conference on Travel Behaviour Research*, Windsor, 2015.

10.   Horni, A., K. Nagel and K.W. Axhausen, eds. *The Multi-Agent Transport Simulation MATSim*. Forthcoming.

11.   Müller, K. and K.W. Axhausen. Using survey calibration and statistical matching to reweight and distribute activity schedules. *Transportation Research Record*, Vol. 2429, 2014, pp. 157-167.

12.   Bundesamt für Statistik. *Mikrozensus Mobilität und Verkehr 2010, Synthesetabellen 2010.* Released 2012. www.bfs.admin.ch/bfs/portal/de/index/news/01.Document.155249.xls. Accessed Apr. 9, 2015.

13.   Dubernet, T. and K.W. Axhausen. Implementing a household joint activity-travel multi-agent simulation tool: First results. *Transportation*, Vol. 42, No. 5, 2015, pp. 753-769.

14.   Balac, M., F. Ciari and K.W. Axhausen. Carsharing demand estimation: Zurich, Switzerland, Area Case Study. *Transportation Research Record*, Vol. 2536, 2015, pp. 10-18.

15.   Bundesamt für Statistik. *Mobilität in der Schweiz – Ergebnisse des Mikrozensus Mobilität und Verkehr 2010*. Released 2012. http://www.portal-stat.admin.ch/mz10/docs/840-1000.pdf. Accessed July 28, 2015.

16.   Santi, P., G. Resta, M. Szell, S. Sobolevsky, S.H. Strogatz and C. Ratti. Quantifying the benefits of vehicle pooling with shareability networks. *Proceedings of the National Academy of Sciences*, Vol. 111, No. 37, 2014, pp. 13290-13294.

17.   Balac, M. and F. Ciari. Enhancement of the carsharing fleet utilization. Presented at *15th Swiss Transport Research Conference*, Ascona, 2015.

18.   Weikl, S. and K. Bogenberger. Relocation Strategies and Algorithms for Free-Floating Car Sharing Systems. *Intelligent Transportation Systems Magazine, IEEE*, Vol. 5, No. 4, 2013, pp. 100-111.

19.   Anderson, J.M., N. Kalra, K.D. Stanley, P. Sorensen, C. Samaras and O.A. Oluwatola. Autonomous Vehicle Technology – A Guide for Policymakers. *Technical Report*, Rand Corporation, Santa Monica, California, 2014.

20.   Boesch, P.M. and F. Ciari. Agent-based Simulation of Autonomous Cars. Presented at *American Control Conference*, Chicago, 2015.

LIST OF FIGURES

FIGURE 1 Start locations of all trips simulated and the network representation used in the MATSim Zurich scenario. The total area shown in the figure is 422.0km x 272.3km.
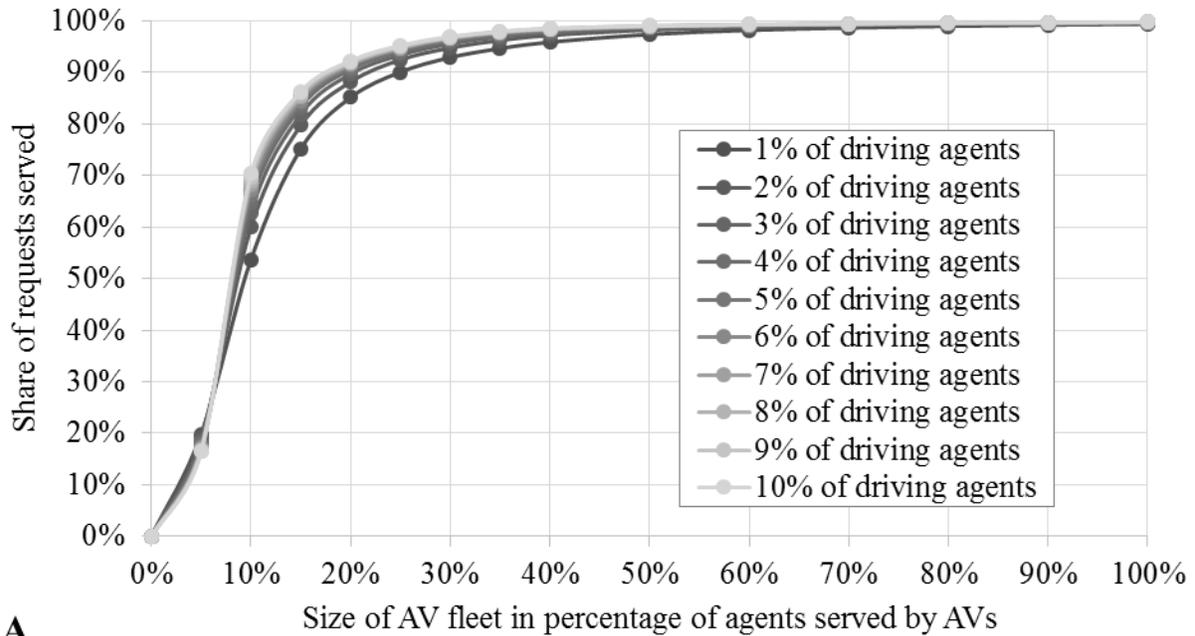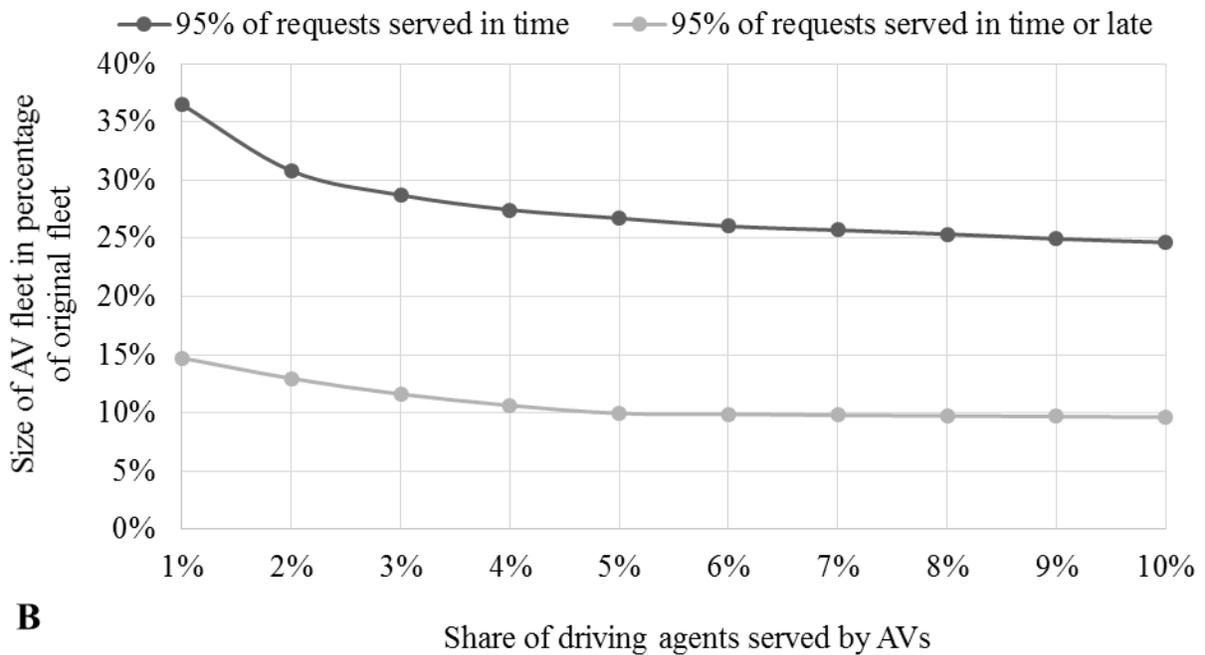
FIGURE 2 AV-Fleet performance versus share of driving agents served by AVs.
Subfigure A shows the performance of differently sized AV fleets for different demand levels.
The performance is measured as the share of requests served within the first 5 minutes after
the requests were registered. Subfigure B presents a crosscut through subfigure A at the
95%-performance level. It shows for a required share of 95% of the requests served within 5
minutes (dark) and within 10 minutes (light) the number of required AVs to serve different
demand levels (1% to 10%).
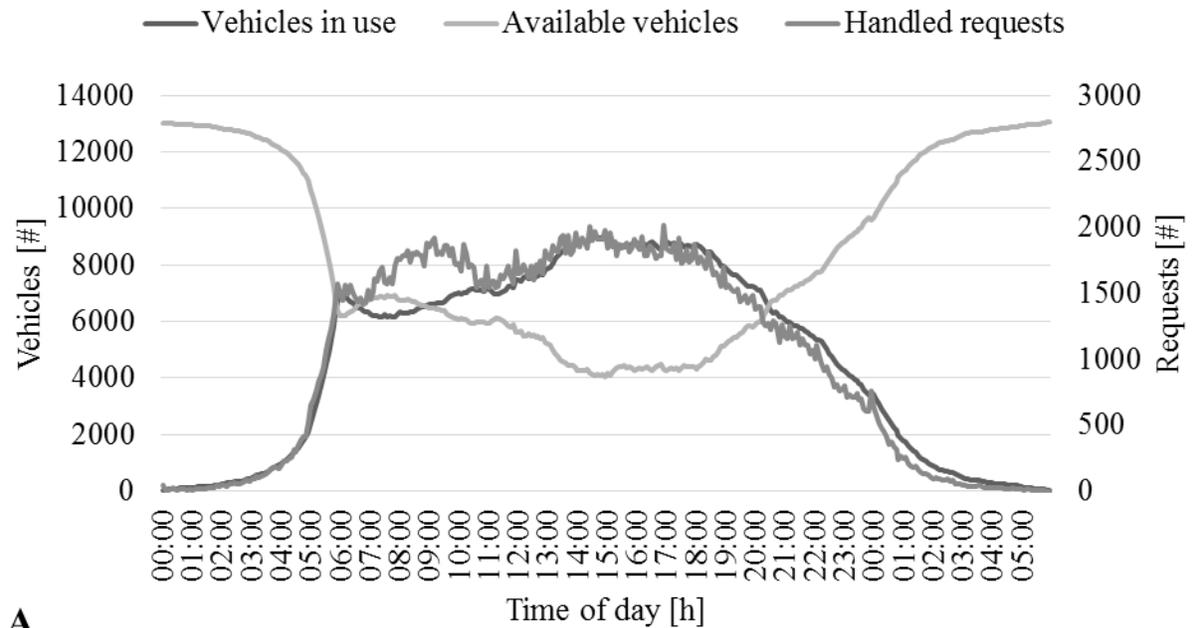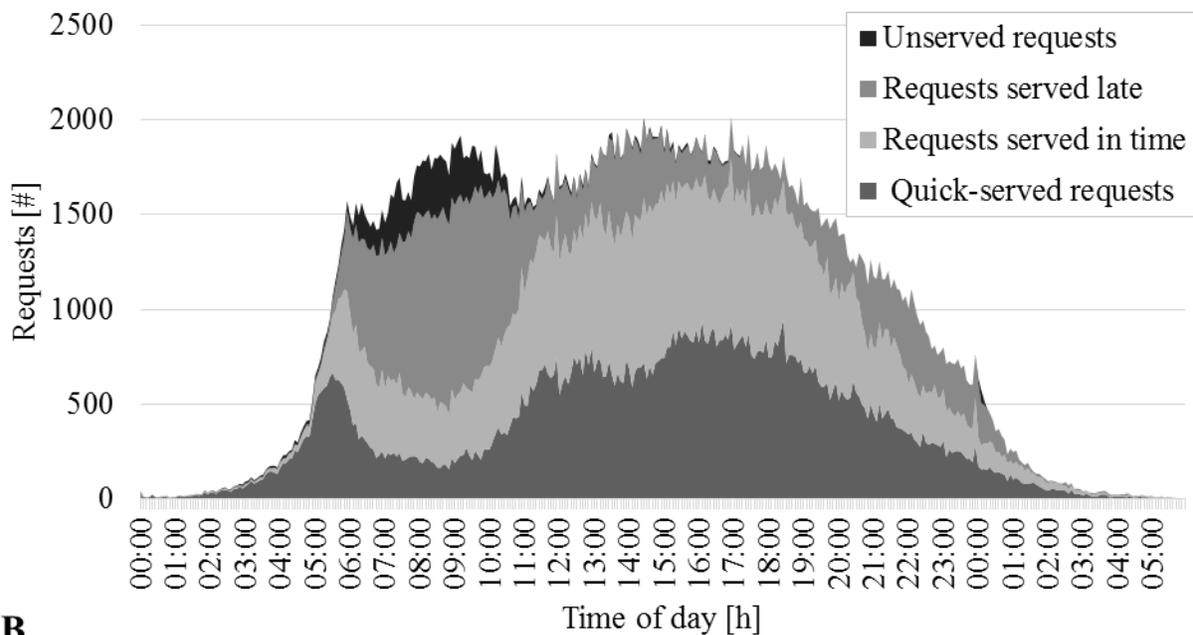
FIGURE 3 Usage and performance of the AV fleet over the course of a simulation day (30h). Subfigure A presents the AVs in use, the available vehicles and the handled requests over the day. The number of vehicles represent snapshots at the time of evaluation. The requests represent all requests served or dropped within the preceding sample interval (5 minutes). Subfigure B shows the number of requests which are quick-served (response within 1 minute), in time (response within 5 minutes), late (response within 5 to 10 minutes) or which could not be served within 10 minutes (unserved requests).
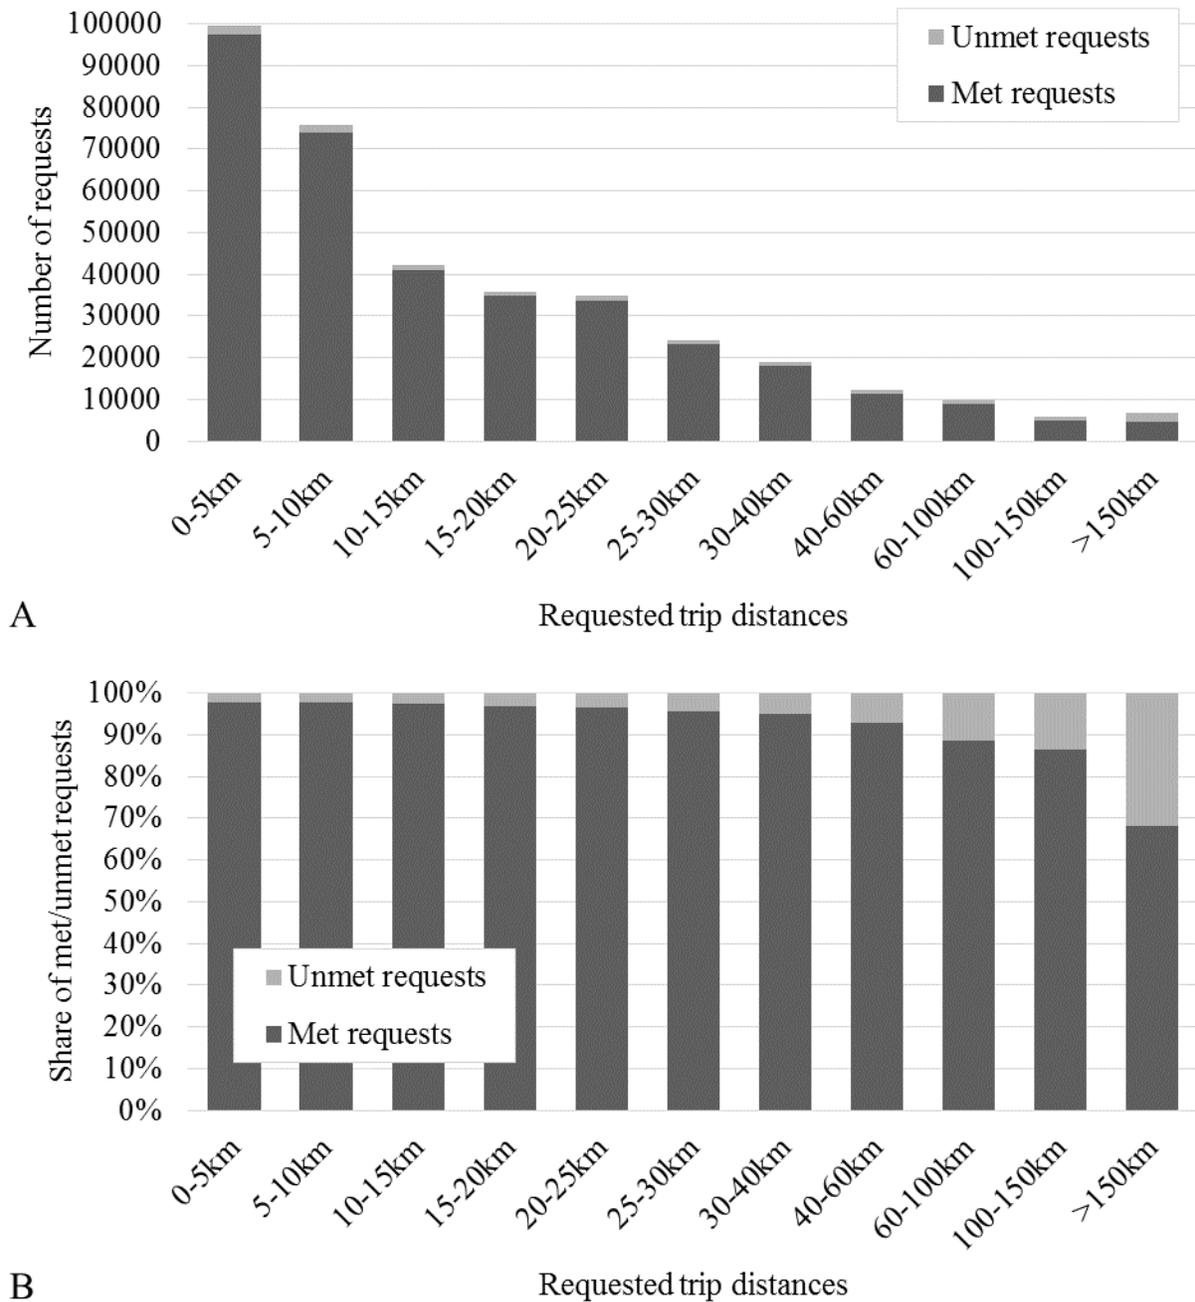
FIGURE 4 Histograms of the requested trip distances.
Subfigure A shows the absolute number of requested trip distances (met within 10min or unmet), subfigure B the relative number of met and unmet requests.
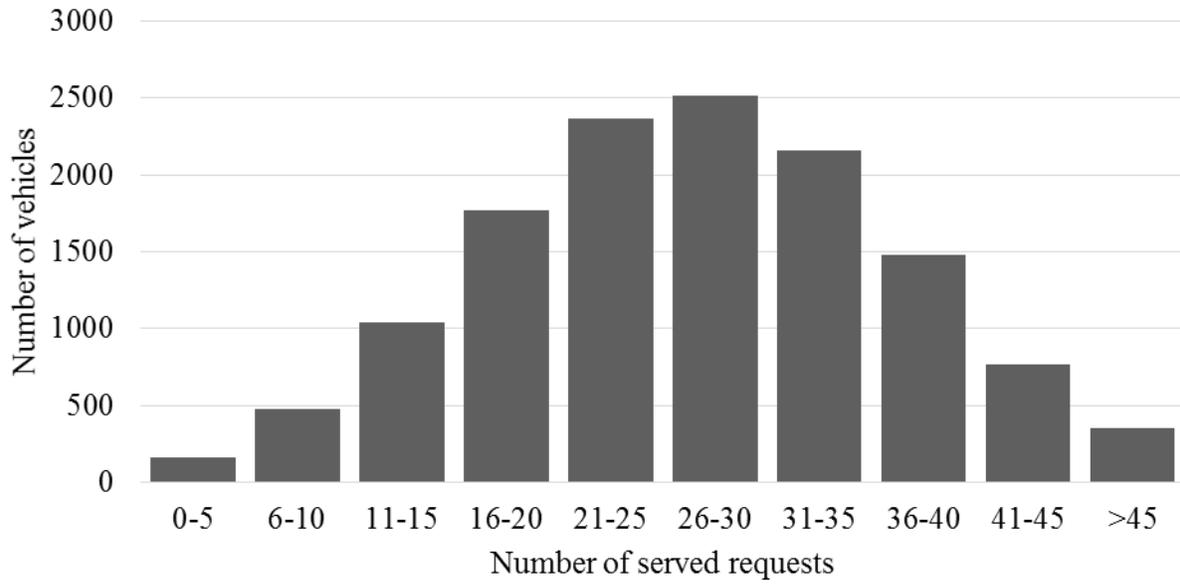
FIGURE 5 Histogram over the full AV fleet showing the number of served requests per vehicle.