



The 6th International Conference on Ambient Systems, Networks and Technologies
(ANT 2015)

Large-scale microscopic simulation of taxi services

Michał Maciejewski^{a,b,*}, Joschka Bischoff^b

^aDivision of Transport Systems, Poznan University of Technology, Piotrowo 3, 60-965 Poznan, Poland

^bDepartment of Transport Systems Planning and Transport Telematics, TU Berlin, Salzufer 17-19, 10587 Berlin, Germany

Abstract

The paper presents research on large-scale microscopic simulation of taxi services in Berlin based on floating car data collected by the Taxi Berlin fleet, the largest taxi association in Germany's capital. Firstly, Berlin's taxi market is shortly described and the demand and supply data obtained from FCD analysed. Secondly, the online taxi dispatching problem formulation for this specific case is given, followed by the definition of two real-time rule-based heuristics used to dispatch taxis dynamically within the simulation. Finally, the simulation setup in MATSim is described, and the results obtained with both heuristics are analysed and compared in terms of dispatching performance, proving the effectiveness of the second strategy at different demand scales.

© 2015 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 Conference Program Chairs

Keywords: online taxi dispatching, dynamic vehicle routing, floating car data, Berlin, MATSim, microscopic simulation

1. Introduction

Simulation is an indispensable tool in order to analyse and optimize complex services with dynamically changing demand and supply, all embedded into a dynamic environment. In the case of taxi services in large cities, a reliable simulation approach must model all these elements both at the microscopic level of detail and in the large-scale (regional, or at least city-wide) scope. The high dynamism of taxi demand is a result of almost all requests being immediate trips with an unspecified destination. For example, partially independent taxi drivers who can choose which rank to wait at, reject serving a request, or decide upon their working hours, account for the limited control over the stochastic supply side. Finally, urban traffic, being the environment for taxi services, implies stochastic time-dependent travel times.

Once the simulation model is ready, analysis and optimization may begin, for instance, by changing the dispatching algorithm, scaling demand and supply, or relocating taxi ranks. As far as the authors know, out of many taxi simulation models^{1,2,3,4,5}, the microscopic ones, though limited in scope, were created for Singapore^{1,4} and Mielec, Poland^{6,7} only. This paper presents a wide-range microscopic model covering the city of Berlin and the neighbour-

* Corresponding author. Tel.: +48 61 647 5957; fax: +48 61 665 2736.

E-mail address: michal.maciejewski@put.poznan.pl, maciejewski@vsp.tu-berlin.de

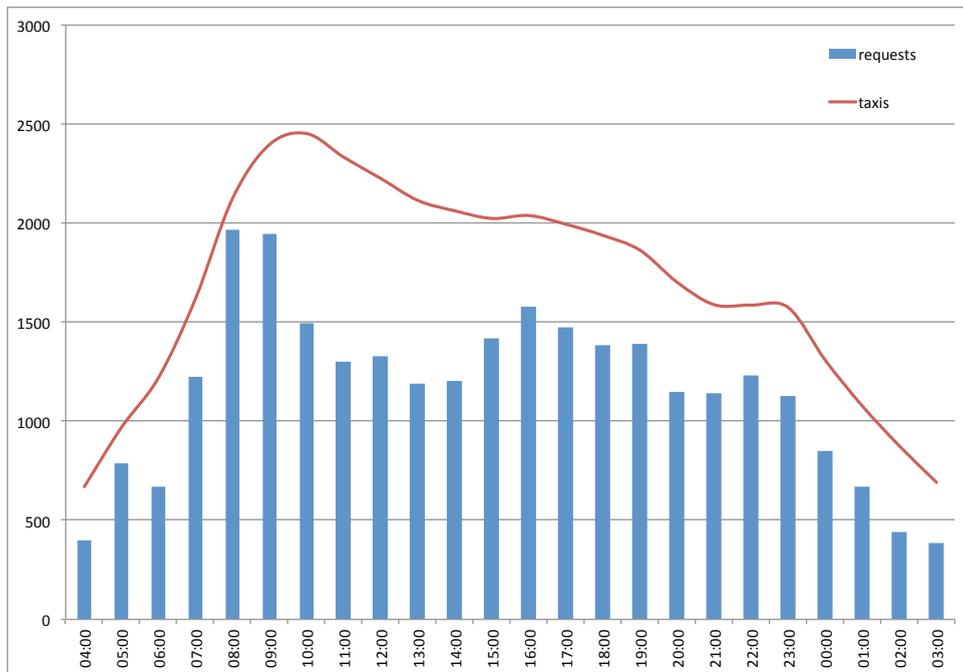


Fig. 1: Hourly request submissions and active taxicabs

ing Brandenburg region, which is then used to assess the performance of two real-time rule-based taxi dispatching strategies.

2. Taxi supply and demand in Berlin

Currently, the Berlin taxi market consists of some 7,600 vehicles licensed to operate in the city. They are organised in 3000 taxi companies employing roughly 18,000 taxi drivers. To model Berlin's taxi supply and demand, in this paper, trajectories of Berlin's biggest radio taxi operator, Taxi Berlin, are used. Overall, they have some 5,700 vehicles working within their range, most of them equipped with GPS trackers that submit their current location and occupation status in a flexible interval (depending on the vehicle's occupation status, but at least once every 60 seconds). These data are, among others, mainly used for travel time prediction in Berlin⁸. With the current occupation status of the vehicle also being submitted, a zone-based matrix of demand for each hour can be generated and used for the demand side of the simulation. On the supply side only the amount of vehicles logged into the system at each second is known, not the actual length of each vehicle or driver shift. This is due to data anonymization by regular reassignment of IDs to vehicles. Furthermore, the amount of vehicles per zone in each occupation status is known in intervals of five minutes.

For the simulation purposes, the supply and demand data of one week (15 April - 22 April 2014) were provided, of which the timeframe between Tuesday 4:00 am and Wednesday 4:00 am has been picked for simulation¹. Figure 1 shows the amount of taxis and requests served during the timeframe investigated. Overall, 27,376 trips were registered. The vehicle supply adjusts to the demand for taxi trips. There is a strong morning peak followed by two somewhat smaller peaks in the afternoon and evening.

The extracted taxi demand is aggregated into 518 zones. Within the city boundaries, these zones correspond to the city quarters defined by the city administration as *Lebensweltlich orientierte Räume (LOR)*⁹. In the outskirts, community boundaries are used instead. The zone with the highest amount of trips starting or ending is the one

¹ Around 4am, taxi supply and demand is the lowest, making it an intuitive breakpoint between days

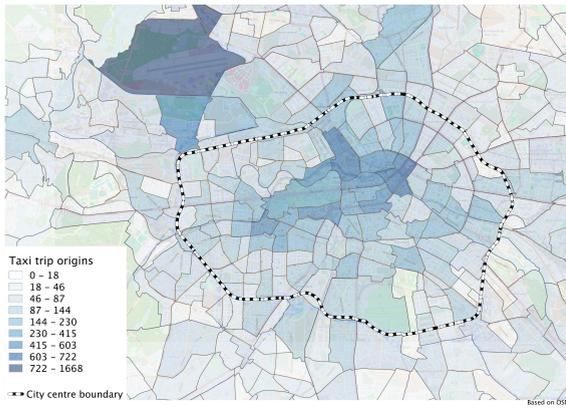


Fig. 2: Spatial distribution of daily taxi demand

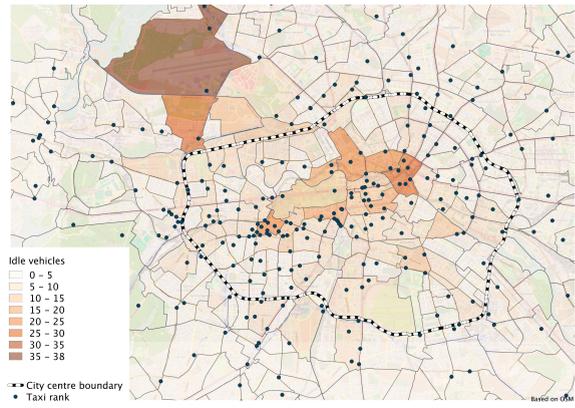


Fig. 3: Average amount of idle taxis per zone

around Berlin's major airport, Tegel. Within 24 hours, 3,799 trips to and from the airport were registered. Most other trips are either ending or beginning in the city centre. Figure 2 shows the origin of taxi trips within the city centre. Idle vehicles tend to aggregate at taxi ranks. Within Berlin, there are roughly 400 of them in operation. Figure 3 provides an overview of rank locations and the average number of idle taxis in each zone. Also in this figure, the area around Tegel airport is clearly notable for the highest amount of idle vehicles. Despite a relatively high demand, waiting times for drivers of several hours are not uncommon there. A detailed FCD analysis is also available¹⁰.

3. Online taxi dispatching

The formulation of the dispatching process was derived from an earlier formulation used for optimizing taxi services in Mielec, Poland^{11,12}. The description below covers immediate requests with unknown destinations, i.e. customers want taxis to arrive as soon as possible and they specify the destination after being picked up, which corresponds to the overwhelming majority of requests in Berlin. However, this type of request is more challenging in terms of online optimization since the provision of the demand data is maximally postponed. As for the supply side, the formulation includes time windows, which in comparison to the original model, add complexity.

Let $N = \{1, \dots, n\}$ be the set of taxi requests (customers). The following sequence of events is related to serving each request $i \in N$ and is illustrated in Figure 4. Taxi customer i calls a taxi (event E_i^{call} , time τ_i^{call}) specifying the pickup location, p_i . Since only immediate requests are considered, the customer's desired departure time is $\tau_i^{\text{dep}} = \tau_i^{\text{call}}$. A selected taxi is dispatched towards p_i at time τ_i^{disp} (event E_i^{disp}), and immediately after arrival, the pickup starts (event E_i^{pick0} , time τ_i^{pick0}). Once the passenger is picked up (event E_i^{pick1} , time τ_i^{pick1}), he or she specifies the destination, d_i , and the taxi sets out immediately. After reaching d_i , the dropoff begins (event E_i^{drop0} , time τ_i^{drop0}). Once the passenger gets out (E_i^{drop1} , time τ_i^{drop1}), the taxi is ready to serve another request. Due to the stochasticity of taxi dispatching, times τ_i^{call} , τ_i^{dep} , τ_i^{disp} , τ_i^{ready} , τ_i^{pick0} , τ_i^{pick1} , τ_i^{drop0} and τ_i^{drop1} are estimated until the respective events take place, and are therefore subject to change.

Request $i \in N$ is *open* if it either has not been planned yet, τ_i^{disp} is not set, or is planned to be served in the future, $\tau_i^{\text{disp}} > \tau^{\text{curr}}$, where τ^{curr} denotes the current time. Let L be the list of all open requests ordered by τ_i^{dep} . Each request i is inserted into L on submission, E_i^{call} , and removed from L on taxi dispatch, E_i^{disp} .

Let $M = \{1, \dots, m\}$ be the set of vehicles. Each vehicle $k \in M$ is available at location o_k within the time window $[a_k, b_k)$. It is assumed that vehicles do not cruise and remain at the dropoff location of the last served customer. When no request has been assigned to k , o_k is k 's initial location and a_k is the time the taxi starts operating. Otherwise, o_k is the dropoff location, d_i , and a_k is the time the dropoff ends, τ_i^{drop1} , of the last request assigned to k , i . Since d_i remains unknown till τ_i^{pick1} , both o_k and a_k are unknown temporarily as well, with the restriction that $a_k > \tau^{\text{curr}}$. Let $M^I \subseteq M$ be the set of all idle vehicles; vehicle $k \in M$ is *idle* if $a_k \leq \tau^{\text{curr}} < b_k$.

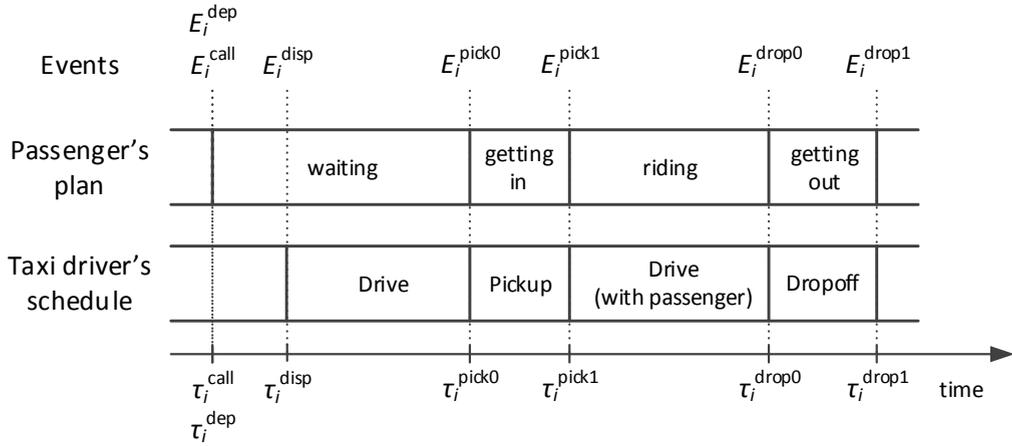


Fig. 4: A taxi driver's schedule and a passenger's plan

4. Rule-based heuristics

Two heuristic dispatching strategies were used to manage the fleet of taxis. Neither strategy uses information about the destinations of the busy vehicles to predict their future availability; thus, the choice of taxis to be dispatched is limited to idle taxis, M^I .

The first heuristic, called *nearest-idle-taxi*, mimics the approach used by many taxi companies. It always serves awaiting requests in the FIFO order by dispatching the nearest idle taxi, $k^* \in M^I$, to the first request in L , $L[1]$. By default, the nearest taxi is defined according to the travel time criterion, that is

$$k^* = \arg \min_{k \in M^I} t_{k,L[1]}^O(\tau^{\text{curr}}), \tag{1}$$

where $t_{ki}^O(t)$, $k \in M$, $i \in N$ is the travel time from o_k to p_i , given the departure time, t .

The strategy reacts to the following events:

- E_i^{call} — if $M^I \neq \emptyset$, vehicle k^* is dispatched to request i
- E_i^{drop1} — if $L \neq \langle \rangle$, vehicle k , after having completed request i , is dispatched to request $L[1]$

This strategy does not look into the future to predict the availability of busy taxis and create schedules². Therefore, even distance-based measures, such as straight-line distance, can be applied to determine k^* . Being very simple, it allows fast, real-time taxi dispatching. However, the main drawback is its poor performance under high demand; when $|M^I| \rightarrow 0$, k^* may be on the opposite side of a city.

To address this problem, a *demand-supply balancing* strategy was used. It classifies the system state into the two following mutually-excluding categories: oversupply ($M^I \neq \emptyset \wedge L = \langle \rangle$) and undersupply ($M^I = \emptyset \wedge L \neq \langle \rangle$) and handles these two situations differently. In the former case, the requests are processed as in the first algorithm, whereas in the latter case, vehicle k is dispatched to the nearest awaiting request, $i^* \in L$. Given the time-based distance measure, k is assigned to serve i^* such that

$$i^* = \arg \min_{i \in L} t_{k,i}^O(\tau^{\text{curr}}). \tag{2}$$

The *demand-supply balancing* strategy reacts to the following events:

² Therefore, in the previous publications, the authors referred to it as the *no-scheduling* strategy (NOS).

- E_i^{call} — if $M^I \neq \emptyset$, vehicle k^* is dispatched to request i
- E_i^{drop1} — if $L \neq \langle \rangle$, vehicle k , after having completed request i , is dispatched to request i^*

The *demand-supply balancing* strategy was based on the assignment rules for scheduling of automated guided vehicles studied by Egbelu and Tanchoco¹³. They proposed order-initiated and vehicle-initiated assignment rules. The former are applied upon request submission (here, E_i^{call}), whereas the latter are executed upon request completion (here, E_i^{drop1}). Under low demand, the *balancing* strategy serves requests immediately as they arrive (request-initiated planning). However, in an overloaded system, the focus is shifted from the FIFO processing order to maximizing vehicle utilization by dispatching vehicles that have become idle to the nearest awaiting requests (vehicle-initiated planning). This results in increased throughput, and consequently, reduces the amount of time passengers await taxis.

5. Simulation setup

MATSim¹⁴ was used as the simulation platform. In general, it allows for running agent-based transport simulations, and its primary purpose is activity-based travel demand modelling. Due to the growing set of extensions, MATSim's core functionality can be extended to serve many purposes.

To simulate taxis, the DVRP (*Dynamic Vehicle Routing Problem*) extension¹⁵ was used. The vehicle routing model implemented in this extension covers a wide range of problem types, offering such features as one-to-one and one-to-many (many-to-one) topologies, multiple depots, dynamic requests, non-homogeneous requests and vehicles, time windows, time-dependent stochastic travel times and costs, network-based routing, and vehicle monitoring and diversion.

The DVRP extension provides the connection between MATSim and the optimization algorithm. It listens to simulation events, monitors the state of simulation, binds driver behaviour to the schedules computed by the optimizer and coordinates interaction between drivers, passengers and dispatchers. It also facilitates optimization by providing the functionality of least-cost path/tree search. For a detailed description of the DVRP extension, the reader is referred to¹⁵.

For the simulation in MATSim, a network based on the OpenStreetMap data is used. It represents the road network in Berlin and the surrounding area of Brandenburg. The network consists of 11,353 nodes and 24,350 links. Using car traffic based on an earlier simulation¹⁶, time-specific link travel times were generated without having to simulate all background traffic. This keeps the simulation runtime reasonable, as only taxi traffic must be considered. The extracted demand data from FCD were converted into plans for MATSim agents. Locations for origins and destinations within the zones were randomly distributed as were the actual departure times within each hour. This resulted in 27,386 agent plans with exactly one taxi trip each. With this demand, the overall stress on the system is low. However, to accommodate a comparably high share of black-market rides in Berlin (30-40% according to estimations)¹⁷, a scaling factor of 1.5 for the taxi demand seems reasonable. Moreover, taxi demand may be much higher due to bad weather conditions, public transport breakdowns, trade fairs or other events. For example, during a recent strike (15 October 2014) in the overground railway system, our data show the amount of taxi trips doubled during the afternoon peak with the taxi supply only increasing by about 20%. On that particular day, a scaling factor for the original demand of 2.5 would be appropriate. To depicture all these fluctuations, the authors decided to scale up the demand step-wise up to 5.0 while keeping the original spatiotemporal distribution of requests.

6. Results

Out of several different performance measures proposed in⁶, the following two are analysed in this section:

- the average passenger waiting time, $T_W = \sum_{i \in N} (\tau_i^{\text{pick0}} - \tau_i^{\text{dep}}) / n$, representing the customers' perspective
- the average pickup trip time, $T_P = \sum_{i \in N} (\tau_i^{\text{pick0}} - \tau_i^{\text{disp}}) / n$, representing the drivers' perspective

Computational experiments were carried out 20 times for each individual setting (i.e. both the strategy and demand scaling). Figure 5 presents the results obtained at different levels of demand. Up to the scaling factor of 2.5 both

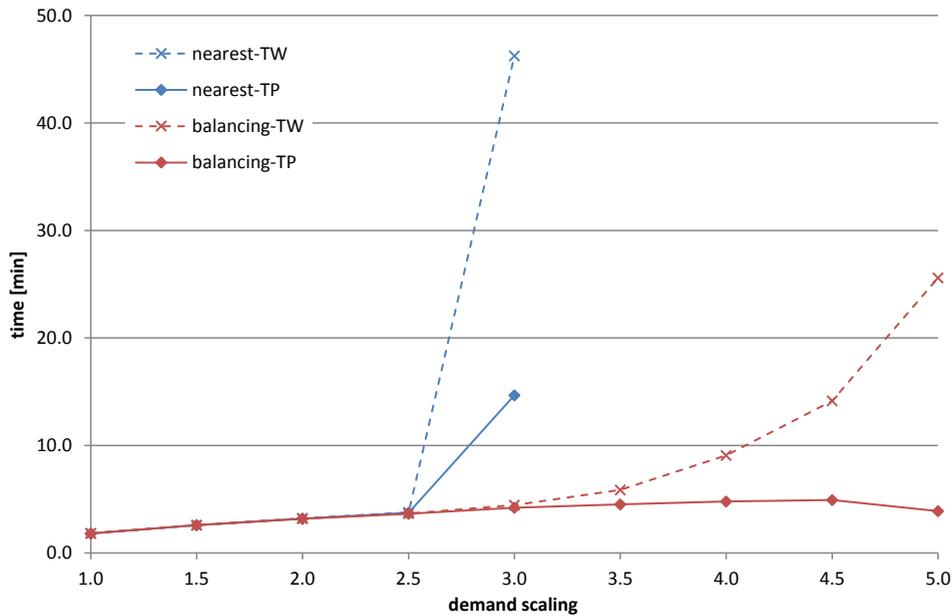


Fig. 5: Average waiting time, T_W , and pickup trip time, T_P , as a function of demand level for the *nearest-idle-taxi* and *demand-supply-balancing* strategies

approaches have comparable performance in terms of T_W and T_P . This is because at low demand only the request-initiated assignments are carried out, i.e. the nearest idle taxi is dispatched to each newly submitted request (the same behaviour in both strategies). However, when the system is overloaded, only vehicle-initiated planning takes place, which leads to significant differences in the performance of both methods, as the *balancing* strategy uses a different rule and tries to minimize the pickup trip times instead of serving requests in the FIFO order. As a result, the *nearest-idle-taxi* strategy fails for the tripled demand (T_W increases to almost 50 minutes), whereas the *balancing* approach efficiently handles even the quadrupled demand (T_W is below 10 minutes).

In general, under high load, the first strategy is extremely myopic as it dispatches vehicles only by looking at the first request in the queue, even though the selected vehicle could be utilized more effectively if dispatched to another request. On the other hand, the other strategy takes into account all awaiting requests and sends the idle vehicle to the closest request which reduces T_P and consequently increases the system throughput. Since $\tau_i^{\text{call}} = \tau_i^{\text{dep}} \leq \tau_i^{\text{disp}}$, $i \in N$, T_P is a component of T_W , thus, achieving relatively small T_P generally leads to smaller T_W . This relation may be observed in Figure 5, where keeping T_P below 5 minutes enables the *balancing* strategy to serve much higher demand. Interestingly, T_P drops for the demand scaling factor of 5.0 and is expected even further as the demand grows. This is caused by the fact that as demand grows, L lengthens, hence the distance to the nearest open request drops on average. Moreover, the *balancing* strategy instantaneously switches between the dispatching rules, even during short interleaving periods of undersupply and oversupply; hence, reducing the risk of getting into larger undersupply.

7. Conclusions

Dispatching a fleet of over 5,000 vehicles in Germany's capital city, Berlin, in order to serve several dozens of thousands of requests daily, poses a challenging optimization problem. The problem is inherently dynamic, as almost all requests are immediate and without pre-specified destinations.

By means of large-scale microscopic simulation run in MATSim, and the dynamic routing functionality offered by the DVRP module, a realistic simulation model of the taxi services in Berlin was developed. This allowed a detailed benchmarking of two rule-based dispatching strategies to be carried out, out of which the second was able to efficiently serve both low and high demand.

Combining both simplicity and efficiency, rule-based dispatching is a very attractive and, thus, overwhelmingly popular approach for management of a fleet of taxis in the real world. However, their apparent disadvantage is the limited planning horizon; the selected decision is the best among all possible single moves, without considering potential sequences of moves. Therefore, the authors plan to use the microscopic taxi model of Berlin to study in detail other, more sophisticated dispatching algorithms, the best of which could be later applied to managing a real fleet.

References

1. Lee, D., Wang, H., Cheu, R., Teo, S.. Taxi dispatch system based on current demands and real-time traffic conditions. *Transportation Research Record: Journal of the Transportation Research Board* 2004;**1882**(-1):193–200.
2. Alshamsi, A., Abdallah, S., Rahwan, I. Multiagent self-organization for a taxi dispatch system. In: *8th International Conference on Autonomous Agents and Multiagent Systems*. 2009, p. 21–28.
3. Wang, H., Lee, D., Cheu, R.. PDPTW based taxi dispatch modeling for booking service. In: *Natural Computation, 2009. ICNC'09. Fifth International Conference on*; vol. 1. IEEE; 2009, p. 242–247.
4. Seow, K., Dang, N., Lee, D.. A collaborative multiagent taxi-dispatch system. *Automation Science and Engineering, IEEE Transactions on* 2010;**7**(3):607–616.
5. Cheng, S., Nguyen, T. Taxisim: A multiagent simulation platform for evaluating taxi fleet operations. In: *Proceedings of the 2011 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology-Volume 02*. IEEE Computer Society; 2011, p. 14–21.
6. Maciejewski, M., Nagel, K.. Simulation and dynamic optimization of taxi services in MATSim. VSP Working Paper; TU Berlin, Transport Systems Planning and Transport Telematics; 2013. See <http://www.vsp.tu-berlin.de/publications>.
7. Maciejewski, M., Nagel, K.. A microscopic simulation approach for optimization of taxi services. In: Albrecht, T., Jaekel, B., Lehnert, M., editors. *Proceedings of the 3rd International Conference on Models and Technologies for Intelligent Transportation Systems 2013*. TUDpress; 2013, p. 1–10.
8. Ebendt, R., Sohr, A., Touko-Tcheumadjeu, L., Wagner, P.. Dynamische Neuplanung der Touren von Express Trucks unter Einbeziehung einer FCD-basierten Verkehrslage. In: *Multikonferenz Wirtschaftsinformatik 2012 : Tagungsband der MKWI 2012*. 2012, p. 124–136.
9. Senate Department for Urban Development and the Environment, . Lebensweltlich orientierte Räume (LOR) in Berlin. 2015. URL: <http://www.stadtentwicklung.berlin.de/index.en.shtml>.
10. Bischoff, J., Maciejewski, M., Sohr, A.. Analysis of Berlin's taxi services by exploring GPS traces; 2015. Accepted for MT-ITS 2015.
11. Maciejewski, M.. Online taxi dispatching via exact offline optimization. *Logistyka* 2014;**3**:2133–2142.
12. Maciejewski, M.. Benchmarking minimum passenger waiting time in online taxi dispatching with exact offline optimization methods. *Archives of Transport* 2014;**30**(2):67–75.
13. Egbelu, P.J., Tanchoco, J.M.. Characterization of automatic guided vehicle dispatching rules. *The International Journal of Production Research* 1984;**22**(3):359–374.
14. Horni, A., Nagel, K., Axhausen, K.W., editors. *The Multi-Agent Transport Simulation MATSim*. (forthcoming); 2015.
15. Maciejewski, M.. Dynamic transport services. In: ¹⁴; 2015, p. forthcoming.
16. Neumann, A.. *A paratransit-inspired evolutionary process for public transit network design*. Ph.D. thesis; Technische Universität Berlin; 2014.
17. Sauer, S.. Wie der mindestlohn das taxi-gewerbe veraendert. *Berliner Zeitung* 2014;URL: <http://www.berliner-zeitung.de/wirtschaft/mindestlohn-2015-in-deutschland-wie-der-mindestlohn-das-taxi-gewerbe-veraendert,10808230,29405460.html>.