MULTI-AGENT BASED SIMULATION OF INDIVIDUAL TRAFFIC IN BERLIN

Kai Nagel
Professor
Group Transportation Systems Planning and Transport Telematics
University of Technology Berlin
Salzufer 17-19
D-10587 Berlin
Germany
Tel: +49 30 314 23308
Fax: +49 30 314 26269
E-mail: nagel@vsp.tu-berlin.de

Arnd Vogel
Researcher
Group Transportation Systems Planning and Transport Telematics
University of Technology Berlin
Salzufer 17-19
D-10587 Berlin
Germany
Tel: +49 30 314 29522
Fax: +49 30 314 26269
E-mail: vogel@vsp.tu-berlin.de

Abstract: Multi-agent simulations of traffic are widely expected to become an important tool for transportation planning in the mid-term future. This paper reports on the first steps of a project which aims to apply such a tool to a large real world scenario based on datasets create in the normal transportation planning process for use with established transportation planning tools. As the first steps of the implementation show, many problems related to different data semantics and the different modelling concepts can occur. In most cases, theses problems can be resolved by minor adoptions of the software or the data. The evaluation of a large scenario of several hundred thousand agents shows that performance issues do no longer hinder applications of this kind. Thereby, the implementation of this scenario helps to push multi-agent traffic simulation tools forward to real world applications.

Keywords: Multi-agent based simulations, Traffic simulations, Activity generation, real world scenarios

MULTI-AGENT BASED SIMULATION OF INDIVIDUAL TRAFFIC IN BERLIN

1 Introduction

In the recent years, microscopic traffic simulations have been an active field of research in transport science. In this context, "Micro" refers to the fact that all elements of the transport systems like roads, crossings and most important, the individual travellers (referred to as "agents") are resolved. This modelling approach is therefore in contrast with the more aggregate models implemented in current transport planning software and used by transportation planners. While these programs have seen several decades of development and practical use, multi-agent based microscopic simulation systems are still relatively new and are mostly used in small and medium scale scientific scenarios rather than in real world applications. But new technologies like high performance computer clusters let multi-agent based simulation systems now become able to run large scenarios (BALMER M. *et a*l ,2004b).

This paper addresses the issue of applying such a model to a real world scenario of large dimensions using input data commonly available for transportation planning. It reports on the first work steps of a research project titled "Environmentally-oriented Road Pricing for Livable Cites – Applications of Agent-based Microsimulation Models to Berlin and Zurich" funded by the Volvo Research and Educational Foundations (VREF), which aims to compare "traditional" aggregate and multi-agent based approaches for the modelling of the potential effects of urban road pricing in Berlin. Some of the essential aspects of this comparative study will be the ability to use the same input data in both approaches and of course the computational tractability of a scenario of this size.

The working plan for this first phase of the project therefore includes the creation of an initial input dataset which can be used by both models based on available data, testing of the software and data integration and evaluation of the computational tractability and performance of the simulation, and the validation of the model output.

2 General simulation structure and special implementation features

In multi-agent based simulation models, traffic is modelled on the level of single individuals – the agents – as opposed to aggregated vehicle counts or traffic streams. The models follow the general approach of demand generation, distribution, mode choice and assignment known from the "classical" 4-step-process of traffic modelling, but the individual steps are replaced by procedures which allow to keep track of the individual agents throughout the entire workflow (BALMER, M. *et al.*, 2004c).

The agents are regarded as a synthetic representation of the real population of the study area. Each agent is assigned one or more plans for one day. These plans are made up of activities (e.g. work, shop,…) and the legs – characterized by their modes and routes - between the activity locations. The activities are bound to specific geographic locations (not zones!) and carry scheduling information (start time, duration) - whereby the initial schedules are taken as “planned schedules” which don’t take into account travel times on the loaded network and which will be adapted within the simulation process.

The creation of the synthetic population and the assignment of activity patterns to the agents can be regarded as the demand generation step of the 4-step-process. The choice of locations for the different activities corresponds to the trip distribution step of the classical approach. Mode Choice can be handled either within the plans generation process by the creation of different sets of plans according to given modal split values or later within the simulation process by allowing the agents to actively choose plans with different modes based on criteria like experienced travel time. Depending on the data being available and on the focus of the study, all these steps can be carried out differently. For example, a synthetic population and the plans can be generated “from scratch” by using census data and time usage patterns, but by making some assumptions, it’s also possible to generate simple populations from existing OD-matrices.

The mobility simulation process, which corresponds to the assignment step of the traditional model, consists of a physical simulation of vehicle movement. Different approaches (e.g. cellular automata, queue models) which model vehicle movement on different levels of detail and which allow for different kinds of evaluations and results can be followed. But besides their strengths, these approaches also differ largely in computational tractability and therefore in the size of populations they can handle. All of them differ from classical assignment procedures in that they are able to generate realistic representations of congestion and their dynamic effects throughout the network.

The mobility simulation step is integrated into an iterative process which allows the agents to evaluate the performance of their plans and to modify different aspects of them (e.g. routes, departure times, mode choice) through so-called “plan mutator modules”. As the plans currently are considered to cover one day, this general workflow of execution – evaluation – modification is called “day to day replanning ” (BALMER, M. *et al.*, 2004a).

The simulation software used in this study is the MAtSim-software (www.matsim.org) which has already been applied successfully to real world studies for the area of Zurich, Switzerland. The unique feature of this software is its ability to operate in parallel on several nodes on a Beowulf computer cluster. Although currently this feature is limited to the mobility simulation module, it allows to simulate large numbers of agents many times faster than real time. Thereby it is possible to run complex scenarios and still do many iterations in reasonable time frames, which is often necessary to allow the system to reach a stable configuration.

3 Study region and input data

The study region covers the urban area of Berlin and the hinterland covering the entire federal state of Brandenburg. The data used in this study is provided by the planning department of the city of Berlin and has used in the major transportation planning activities for the city, namely the creation of the "Stadtentwicklungsplan Verkehr" (SENATSVERWALTUNG FÜR STADTENTWICKLUNG, 2003) since 1998. Due to this origin, the urban area is represented with a much higher level of detail and accuracy than the hinterland in the datasets. The study region covers an area of 150 x 250 km and has a population of about 6 million people.

The data used in this study consists of the "official" road network used in transportation planning which is provided by the planning department, the "Senatsverwaltung für Stadtentwicklung, Bauen und Wohnen" in the VISUM-file format (PTV AG, 2005). It represents only the major road network and consists of about 13000 links.

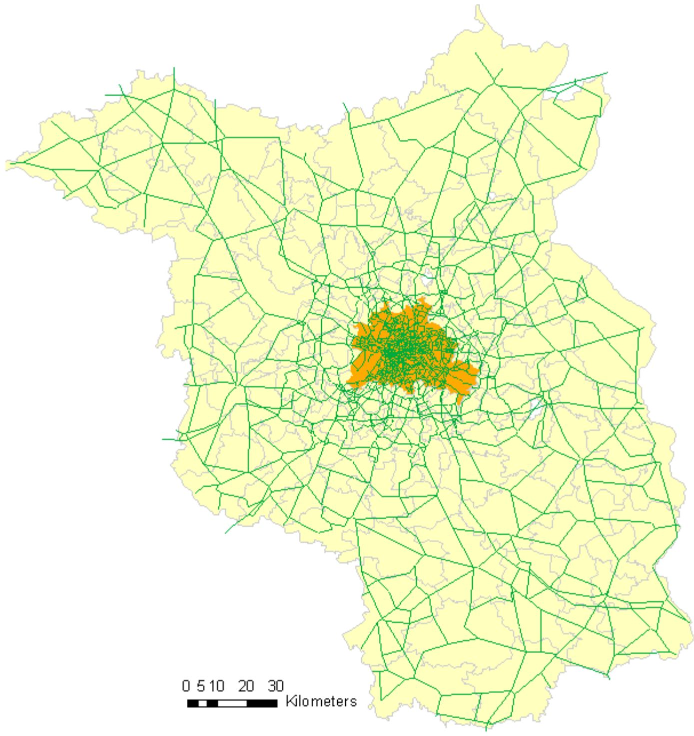

Figure 1: Study region and road network

The demand data is provided in 24h OD-matrices which are based on the 1020 transportation planning zones of the study region. These matrices have been generated by a proprietary transportation model which has been developed and configured specifically for the needs of Berlin (KUTTER, 1990). These OD-matrices represent a total demand of 6.6 million trips per day. No further level of detail like time-dependent matrices or matrices separated by trip purpose is available.

4 Generation of a synthetic population and day plans

As stated above, the demand data which is available for Berlin does only consist of 24h OD-matrices. This kind of demand data is usually sufficient for mid term or long term transportation planning. But multi-agent based simulations require the demand to be given on the individual level as one-day-plans containing the activities with the scheduling and location information and the trips between them.

Several approaches with increasing level of complexity can be taken to generate these more detailed datasets. The simplest approach is to directly convert the trips contained in the OD-matrix to plans with simple artificial activity patterns such as “home-work-home”. Some assumptions about the actual activity locations and the scheduling have to be made in order to fill in the missing information required by the micro-simulation. The simplest approach is to use zone centroids as the activity locations and to have all plans initially start at the same time and have the same duration. One has to account for the fact that the 24h OD-matrix does contain both, the trips leading to the activities and the return trips which can very roughly be done by using only half of the trips. Although it is obvious that an approach like this cannot produce any meaningful results in terms of the real traffic situation, it is useful to generate the plans for initial tests of the road network, because the volume and spatial distribution of the trips are at least roughly related to the network layout. The major drawback is the missing scheduling information which is crucial for the dynamic mobility simulation. Due to the simple assumptions being made, all trips from the 24h-OD-matrix will initially start at the same time leading to high loads in the first iterations which are then relaxed by the replanning modules as described in section 2.

A more realistic approach for the micro-demand-generation is to use demographic data available for the zones in order to account for different trip purposes. The approach is still based on simple activity chains with only one out-of-home activity. The demographic data like age groups, employment rates and number of working places can be used to extract some trip purposes or “primary activities” which have relatively fixed schedules like “work”, “school”, “child care”. This also helps to better sort out the return trips. The remaining trips from the OD-matrix can be arbitrarily distributed across the day. The demand generated by this approach should be realistic enough at least to compare the results for the peak hours to real traffic counts in order to adapt the model parameters.

The third and most advanced approach is to generate the plans of the agents solely based on demographical and landuse data and on activity patterns derived from census surveys such as the time budget / time use census “Zeitbudgeterhebung” (STATISTISCHES BUNDESAMT DEUTSCHLAND, 2005) or the mobility-related “Mobility in Germany” MiD (BUNDESMINISTERIUM FÜR VERKEHR, BAU- UND WOHNUNGSWESEN, 2002) . In this process, a synthetic population is created based on microgeographical demographic data containing information about age groups, employment, household structure and car availability. The individuals of this population are classified into certain demographic groups and then assigned activity patterns with initial arbitrary schedules based on the census data. The activity locations are chosen in two steps. First, a location for primary activities such as “work” or “education” is chosen corresponding to the landuse and working place data. Secondary activities like “leisure” or “shopping” are then chosen by taking into account the home and primary activity locations.

In all cases, the routes between the activity locations are initially computed based on the travel times defined by free speed on the links. During the simulation process, new routes are generated for a portion of the agents based on actual travel times on the loaded network. Furthermore, the iterative simulation process allows the agents to adopt their travel patterns in terms of departure times and to some extend also activity durations in order to optimize the performance of the day plans (BALMER M. *et.al* 2004c).

5 Model setup and necessary adoptions

In the model setup chosen for this initial phase of the study, the demand was generated following the first and most simple approach, because the main goal of this working step is to get the software working and to check whether the provided network data can be used immediately.

In order to speed up the initial test, the demand and the network capacities where scaled down to 10% of the actual values. This still yields a number of more than 150.000 individual agents which have to be modelled.

Although the MAtSim software has already been successfully applied to large scenarios, e.g. for the Kanton Zurich (BALMER, M. *et al*., 2005), the use of an externally prepared network and the very naïve approach to demand generation implied by the simple first approach described in section 4 made many problems appear very quickly in the initial model runs.

The first set of problems relates to the way in which some aspects of the VISUM-road network are coded. Oneway roads are in many cases not simply locked for cars (as is assumed by the converter software of the MAtSim-package) but simply coded with zero values for capacity, free speed, travel time and length. In other cases, capacity values or lengths assigned to links are obviously too low due to coding errors.

In aggregate models such as VISUM, network inconsistencies like these are either tolerated or they lead to small distortions of the results. In the dynamic “physical” mobility simulation of MAtSim, these inconsistencies led to more or less complete stalls of the vehicle flow! The invalid oneway-links with travel time zero which are mistakenly included in the network by the converter software are “very attractive” to the router module in the initial route generation based on free travel times. Therefore, huge volumes are loaded onto these links – and even in the wrong direction of traffic flow. To solve this problem, the converter software had to be adapted to filter out inconsistent links

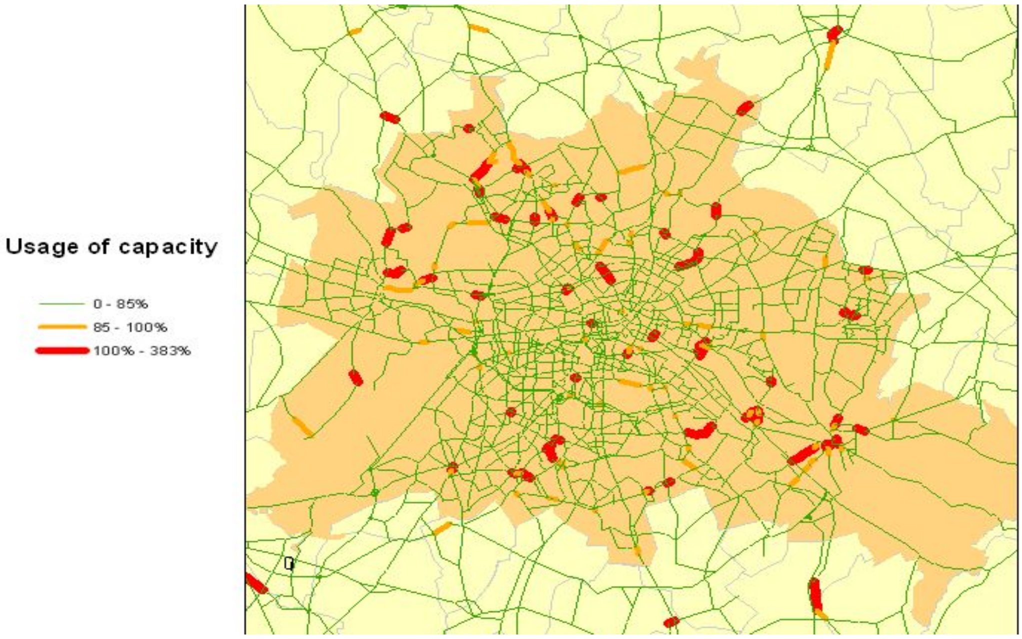

Figure 2: Usage of capacity in VISUM-Assignment

Another problem was caused by the coding errors of short links or small capacities. Figure 2 shows how the network inconsistencies lead to high usage of capacity in a VISUM-assignment. From the spatial distribution of these links – which are widespread across the network and show no clustering in certain areas – it is clear that they are likely caused by network preparation errors rather than the traffic situation.

Again, inconsistencies like these do only distort the results of aggregate models but can lead to more or less complete stalls in the dynamic mobility simulation of MAtSim. "Complete stalls" in this case means that some agents require very long time (several days) for their plans. This is due to the fact that capacities are interpreted very differently by the two models. In aggregate models such as VISUM, the capacity of a link is a value which relates the actual speed on the link to the volume by a capacity constraint function such as BPR. In a "physical" mobility simulation, the capacity as well as the link length is treated as hard constraints which limit the number of vehicles allowed to be on the link at the same time. Furthermore, the transition of vehicles from one link to another at the network nodes is controlled by a stochastic function which calculates transition probabilities dependent on the free slots on the next link. The faulty link lengths and capacities yield to very low transition probabilities even for empty links and therefore completely stall traffic flow on that link. The resolution to this problem was the introduction of a more deterministic transition function which lets vehicles easier transfer onto empty links.

The above mentioned problems where even enhanced by the naïve approach employed for demand generation. The replacement of true activity locations with zone centroids and the initial scheduling which let all agents depart at the same time (7 am) increased traffic volumes on certain links dramatically.

Altogether, these problems resulted in initial simulation periods of up to 200 hours for plans which only included “start at 7am, go to work, work for 8 hours and return home”.

The long running plans finally revealed another problem with of the simulation software. The software is meant to work on day plans and does therefore know nothing about weekdays. The long running plans could thereby result in agents returning earlier than they started, yielding up in negative plan durations. The negative durations can not be evaluated in a meaningful manner by the logarithmic scoring function which is used to select the plans with the best performance. Observable effects of these problems are for instance agents staying at work for many hours. For this problem, only workarounds have been introduced up to now, a true solution is not obvious and would certainly require major changes to the scoring model.

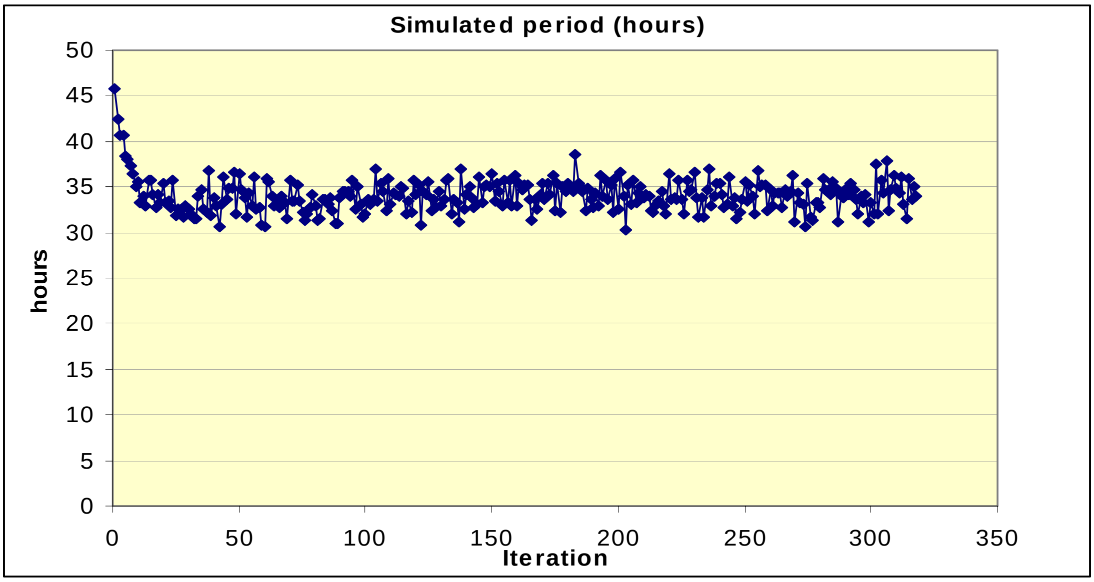

Figure 3: Period simulated by one iteration

The adoptions mentioned above resulted in a working simulation model which operates on a large scenario at a high performance. But these adoptions did not yet result in results which appear to be realistic, because even after more then 300 iterations, we still get agents who have plan which require more than 30 hours to execute. As Figure 3 shows, the overall period simulated in one iteration – the time between the first departure of an agent and the last return of an (usually not the same!) agent drops drastically in the first 20 iterations but then fluctuates around 33 hours – which is still much longer than a day and must therefore be considered problematic. The rapid drop in the first iteration mainly results from changes in route choice of the agents. In the first iteration, routes are chosen based on free speed on the links, causing many agents to chose the same routes – and therefore ending up in congestion. In the next couple of iterations, the dynamic router distributes the route choices among different alternatives, leading to better balanced network loads, less congestion and overall quicker travel. But these effects become quite stable after some iterations and no significant further improvement is seen thereafter.

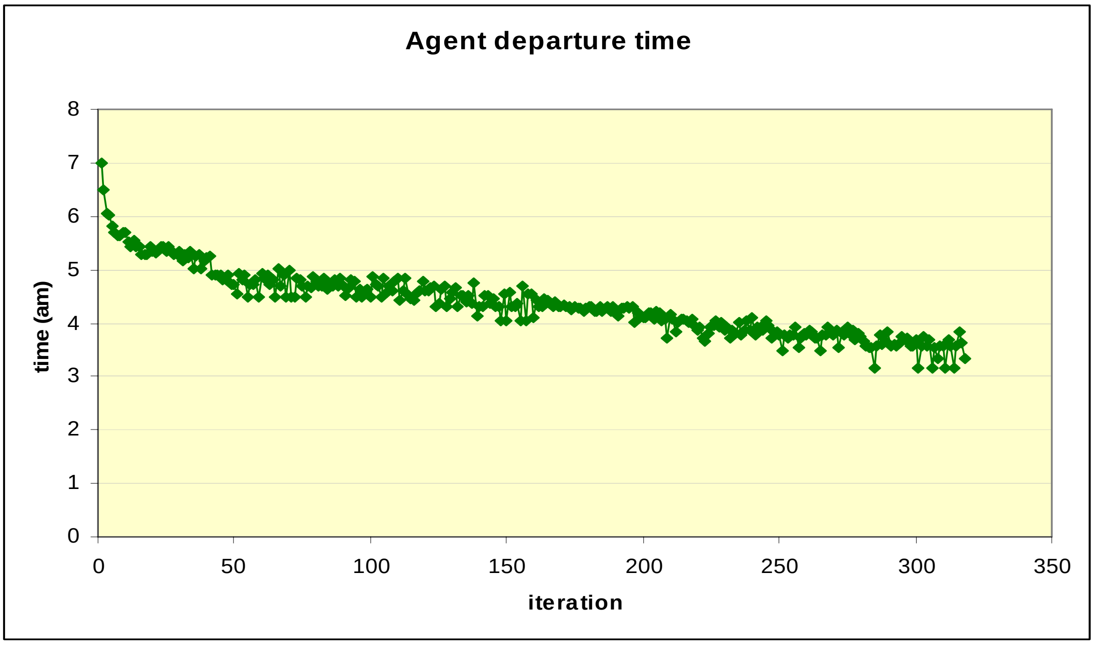

Figure 4: Departure times of the agents (earliest departure)

In general, the strategies to optimize plans performance found by the agents through the iterative learning process show a strong tendency towards modification of departure times. Figure 4 shows how the departure times of the agents get earlier with almost every iteration. Detailed investigation of the results revealed, this behaviour and the long simulation periods mentioned above may be caused by the allocation of agent home and activity locations to links. The simple approach used for this assignment (see section 4) causes many agents to have the same start and destination locations. This in turn leads to long-lasting congestion on the links affected. A solution to this will be the application of a smarter spatial distribution of the locations.

6 Computational tractability and performance

Some of the computational properties of the MAtSim software package are already known from other study which used variants of this software (BALMER M. *et al*., 2005 / BALMER M *et al.* 2004b). Those and earlier studies show that computer memory is the main limiting resource to the number of agents in the simulation. This is even the case although the MAtSim-software is able to use several computers in parallel. The reason for this is that currently, only the mobility simulation module of MAtSim is able to make use of a parallel Cluster architecture by splitting the network geographically into separate pieces which are handled by different node. The other modules are still required to run only one single node, which is due to the fact how the agents and plans information is currently handled.

The agents are managed in an “Agent Database” which is kept in memory during the entire simulation run. Every agent keeps multiple plans – in our study we use 6 plans per agent – in order to have alternatives to chose from. Furthermore, the mutator modules need some memory during their runs. Our test runs, which used a number of up to 150.000 individual agents produced an in-memory agent database of approx. 1 GB of the available 2 GB of RAM. But because of the requirements of the other modules and the mobility simulation, it will be difficult to increase the number of agents to a very much higher number.

As a consequence of this (expected) situation, a new version of the Agent Database is currently being developed. This Agent database will be able to run in parallel mode with several instances of the database, each managing a portion of the entire agent-population. This approach will also allow to run all the mutator modules in parallel mode.

On the hardware used for this study, which consists of common PC hardware (64-Bit AMD Opteron processors at 2.2 GHZ, 2 GB RAM), the software is able to run an iteration using the synthetic population of approx. 150.000 agents in about 23 minutes in Single-CPU-mode (all computation including mobility simulation takes place on only one node). Therefore, the system currently completes 62 iterations per day or is simulating at 62 times real time speed. This performance is already sufficient to run complex scenarios which require several hundred iterations in reasonable time, but further performance enhancements are possible by running the simulation in parallel mode.

7 Conclusion

The implementation of a multi-agent based traffic simulation based on data taken from planning processes which are based on established planning tools is possible, but many traps are on the way to a working simulation system. But these problems expose bugs and conceptual difficulties in the multi-agent simulation software which are very hard to find by using only artificial test datasets. Therefore, the implementation of a real world scenario is very helpful to improve the simulation software.

Using the currently available hardware, the implementation of large scenarios is no longer a problem. The use of sophisticated software designs such as parallel computing will allow to analyze scenarios of any size, opening the way for new research questions which examine the effects of the behaviour of large populations based on decisions taken at the individual level.

ACKNOWLEDGEMENTS

We would like to thank the Volvo Research and Educational Foundations for granting funding for the research project “Environmentally-oriented Road Pricing for Livable Cites”, for which this study serves as preparationary work. Also, we’d like to thank the Senatsverwaltung für Stadtentwicklung, Bauen und Wohnen and the Statistisches Landesamt Berlin for providing the data used in this study.

REFERENCES

M. Balmer, N. Cetin, K. Nagel, B. Raney (2004a) Towards truly agent-based traffic and mobility simulations, **Autonomous agents and multiagent systems** (AAMAS'04), New York, NY. 2004

M. Balmer, B. Raney, K. Nagel (2004b) Large-scale multi-agent simulations for transportation applications**, ITS Journal**, in press

M. Balmer, K. Nagel, B. Raney (2004c) Coupling activity-based demand generation to a truly agent-based traffic simulation - activity time allocation, **EIRASS workshop on Progress in activity-based analysis**, Maastricht, NL, May 2004

Balmer, M., M. Rieser, A. Vogel, K.W. Axhausen and K. Nagel (2005) Generating Day Plans Based on Origin-Destination Matrices. A comparison between VISUM and MATSIM based on Kanton Zurich data. **Proceedings of Swiss Transport Research Conference (STRC)**, Monte Verita, Switzerland, 9-11 March 2005

Bundesministerium für Verkehr, Bau- und Wohnungswesen (2002) Mobilität in Deutschland **http://www.kontiv2002.de/,** accessed April 2005

Kutter, E., H.-J. Mikota, (1990) , Weiterentwicklung des Personenverkehrs-modells Berlin, **Arbeitsberichte zur Integrierten Verkehrsplanung**, University of Technology Berlin (Typoskript).

PTV AG (2005) VISUM Product Information, **http://www.ptv.de/cgi-bin/traffic/traf_visum.pl**, accessed April 2005

Senatsverwaltung für Stadtentwicklung Berlin (2003) Mobil 2010 – Stadtentwicklungsplan Verkehr Berlin, **http://www.stadtentwicklung.berlin.de/planen/stadtentwicklungsplanung/de/verkehr/download/Mobil2010.zip**, accessed April 2005

Statistisches Bundesamt Deutschland (2005) Zeitbudgeterhebung, **http://www.destatis.de/micro/d/micro_c4.htm**, accessed April 2005