Improving a large-scale agent-based simulation scenario

2008-Jul-31

Yu Chen
Transport Systems Planning and Transport Telematics,
Berlin Institute of Technology, D-10587 Berlin
phone: +49-30-314 78783
fax: +49-30-314 26269
chen@vsp.tu-berlin.de

Marcel Rieser
Transport Systems Planning and Transport Telematics,
Berlin Institute of Technology, D-10587 Berlin,
and
Institute for Transport Planning and Systems, ETH Zurich, CH-8093 Zurich
phone: +49-30-314 25258
fax: +49-30-314 26269
rieser@vsp.tu-berlin.de, rieser@ivt.baug.ethz.ch

Dominik Grether
Transport Systems Planning and Transport Telematics,
Berlin Institute of Technology, D-10587 Berlin
phone: +49-30-314 29521
fax: +49-30-314 26269
grether@vsp.tu-berlin.de

Kai Nagel
Transport Systems Planning and Transport Telematics,
Berlin Institute of Technology, D-10587 Berlin
phone: +49-30-314 23308
fax: +49-30-314 26269
nagel@vsp.tu-berlin.de

Words: 6000 + 6 figures + 2 table

ABSTRACT

All real world transport modelling exercises consume considerable resources for the generation of a base case. This is not different for microscopic, behavior-based (or agent-based) models. This paper reports some of the measures that were taken to improve a fully microscopic model of traffic in the Zurich region in Switzerland. Improvements to the network were possible when considering data from the open source project "OpenStreetMap". Other improvements were achieved by adding choice dimensions: besides the usual route choice, the simulated travelers can adjust their activity timing, and their mode choice. Addition of both choice dimensions improved the results, in spite of the already relatively good initial data availability for the Zurich region. We conclude that public data sources will eventually remove some of the data problems for large scale systems, and that the additional adaptive capability of microscopic, behavior-based models may help to make the models more realistic, without much additional calibration burden.

INTRODUCTION

Many cities or regions invest considerable resources into their transport modelling. There are many examples; some of the ones with publications in the academic world are San Francisco *(1)*, Portland *(2)*, New York *(3)*, Chicago *(4)*, Eastern Denmark *(5)*, or Switzerland *(6)*. In all these cases, it seems that either considerable resources are necessary, or the model building process proceeds over many years, with the corresponding experience accumulating incrementally.

The situation is no different for microscopic, behavior-based (or agent-based) transport models: Similar amounts of work are necessary to obtain access to data sources, to merge these data sources, and to gain experience with the strengths and weaknesses of the data sets, in particular vis-a-vis the models. This paper reports the results of such an excercise undertaken for the metropolitan region of Zurich in Switzerland. While Refs. *(7)* and *(8)* describe the demand generation and report some initial results for the whole of Switzerland, this paper concentrates on the region of Zurich and on making the traffic flows more realistic.

Three different types of modifications will be considered:

- Adaptations of the network, where the open source project openstreetmap (www.openstreetmap.org) turned out to be very helpful;
- adaptations of the demand, where the inclusion of long-distance traffic turned out to be beneficial;
- and finally, and most importantly, integrating time adaptation and mode choice as additional adaptive choice dimensions. Thus, these are no longer fixated in the upstream demand generation, but are adapted in an iterative procedure in the same way routes are adapted in dynamic traffic assignment.

It is, in our view, a very positive effect that making additional choice dimensions adaptive makes the base case more realistic. Presumably, the adaptive agents find better ways to adjust to the particularities of the system than the more aggregated upstream methods. This was in spite of the comparatively good initial data availability for the Zurich region.

What is quite different between the approach described here and many other approaches including those mentioned earlier is that at this point the approach described here does not formally calibrate parameters (as could for example be done with BIOGEME *(9)*. Instead, parameters are usually set to plausible values, and then emergent properties of the model (such as hourly traffic flows) are compared to real world data. This has to do with the fact that calibrating agent properties based on simulation-based emergent effects is not straightforward. Also, it is so far our experience that insight into the model behavior is also a successful strategy to build a more realistic model. Nevertheless, Refs. *(10, 11)* indicate that it is possible to develop concepts to calibrate agent-based travel behavior models. This will be the subject of future work.

This paper is structured as follows. First, the simulation structure is explained. This is, except for the mode choice, similar to earlier expositions of the same material. Then, the scenario setup is reported, which contains a short summary of the demand generation process and lists the scenario-specific simulation parameters. Next, the validation methodology is presented, which essentially consists of time-dependent relative error when compared to real world counting stations. A longer section on "improvements" discusses the three elements mentioned above: network modifications, demand modifications, and additional choice dimensions. The paper is finished with a discussion and a conclusion.

SIMULATION STRUCTURE

Our simulation is constructed around the notion of agents that make independent decisions about their actions. Each traveler of the real system is modeled as an individual agent in our simulation. The overall approach consists of three important pieces:

- Each agent independently generates a so-called *plan*, which encodes its intentions during a certain time period, typically a day.
- All agents' plans are simultaneously executed in the simulation of the physical system. This is also called the *traffic flow simulation* or *mobility simulation*.
- There is a mechanism that allows agents to *learn*. In our implementation, the system iterates between plans generation and traffic flow simulation. The system remembers several plans per agent, and scores the performance of each plan. Agents normally chose the plan with the highest score, sometimes re-evaluate plans with bad scores, and sometimes obtain new plans by modifying copies of existing plans.

The simulation approach is the same as in many of our previous papers *(*e.g. *12)*. The following exposition is a shortened and simplified description of key elements to limit the length of this paper. The results of this paper are based on a re-implementation of the MATSim framework in Java *(13)*.

A **plan** contains the itinerary of activities the agent wants to perform during the day, plus the intervening trip legs the agent must take to travel between activities. An agent's plan details the order, type, location, duration and other time constraints of each activity, and the mode, route and expected departure and travel times of each leg.

A plan can be modified by various **modules**. This paper will make use of the following modules:

- *Activity Times Generator:* This module is called to change the timing of an agent's plan. At this point, a simple approach is used which applies a random "mutation" to the duration attributes of the agent's activities. Although this approach is not very sophisticated, it is sufficient in order to obtain useful results. This is consistent with our overall assumption that, to a certain extent, simple modules can be used in conjunction with a large number of learning iterations *(*e.g. *14)*.
- *Router:* The router is implemented as a time-dependent Dijkstra algorithm. It calculates link travel times from the output of the traffic flow simulation. The link travel times are encoded in 15 minute time bins, so they can be used as the weights of the links in the network graph. The algorithm is accelerated using a Landmark A* preprocessing *(15)*.
- *Mode choice* will be simulated by giving every agent both a "car" and a "non-car" plan.

The **traffic flow simulation** executes all agents' plans simultaneously on the network, and provides output describing what happened to each individual agent during the execution of its plan. The traffic flow simulation is implemented as a queue simulation, which means that each street (link) is represented as a FIFO (first-in first-out) queue with two restrictions *(16, 17)*. First, each agent has to remain for a certain time on the link, corresponding to the free speed travel time. Second, a link storage capacity is defined which limits the number of agents on the link. If it is filled up, no more agents can enter this link.

The outcome of the traffic flow simulation (e.g. congestion) depends on the planning decisions made by the decision-making modules. However, those modules can base their decisions on the output of the traffic flow simulation (e.g. knowledge of congestion) using **feedback** from the multi-agent simulation structure *(18, 19)*. This sets up an iteration cycle which runs the traffic flow simulation with specific plans for the agents, then uses the planning modules to update the plans (**replanning**). These changed plans are again fed into the traffic flow simulation, etc, until consistency between modules is reached.

The feedback cycle is controlled by the **agent database**, which also keeps track of multiple plans generated by each agent, allowing agents to reuse those plans at will. The repetition of the iteration cycle coupled with the agent database enables the agents to learn how to improve their plans over many iterations. This circle continues until the system has reached a relaxed state. At this point, there is no quantitative measure of when the system is "relaxed"; we just allow the cycle to continue until the outcome seems stable.

In order to compare plans, it is necessary to assign a quantitative **score** to the performance of each plan. In principle, arbitrary scoring schemes can be used *(*e.g. prospect theory *20)*. In this work, a simple utility-based approach is used. The elements of our approach are as follows:

- The total score of a plan is computed as the sum of individual contributions consisting of positive contributions for performing an activity and negative contributions for travelling and for leaving an activity early.
- A logarithmic form is used for the positive utility earned by performing an activity:

  $$U_{perf,i}(t_{perf,i}) = \beta_{perf} \cdot t_{*,i} \cdot \ln\left(\frac{t_{perf,i}}{t_{0,i}}\right)$$

  where t_{perf} is the actual performed duration of the activity, t_* is the "typical" duration of an activity, and β_{perf} is the marginal utility of an activity at its typical duration. β_{perf} is the same for all activities, since in equilibrium all activities at their typical duration need to have the same marginal utility.
  $t_{0,i}$ shifts the curve vertically and has, as long as activities cannot be dropped, little effect.[1]
- The (dis)utility of traveling is assumed as linear in the travel time with different valuations of the time for two transport modes.
- The (dis)utility of leaving an activity early is assumed as linear in the reduction of the duration of this activity compared to its "typical" duration.

It is important to note that the score thus takes into account the complete daily plan. More details can be found in *(21, 12)*.

SCENARIO SETUP

The network initially used is a Swiss regional planning network *(22)* which includes the major European transit corridors and covers the area of Switzerland. It consists of 24 180 nodes and 60 492 links with attributes (flow capacity, free speed, number of lanes, …) suitable for static traffic assignment, but not for our dynamic agent-based simulation.

[1]There is the well-intentioned rule implemented in MATSim that activities whose score according to the above rule becomes negative should receive score zero, since an agent could always do "nothing" instead. Unfortunately, this means that $t_{0,i}$ now has an effect, since it determines at which duration this rule kicks in. With the simulations reported here, the time pressure should never be so large that this effect is triggered.

An initial demand was prepared that consists of all travelers within Switzerland. The demand generation process is described in more details in Refs. *(7, 8)*. The following paragraphs give a short summary of the most important points for better understanding of the following improvements.

All travelers have complete daily activity patterns based on microcensus *(23)* information. Such activity patterns can include activities of type *home*, *work*, *education*, *shopping*, and *leisure*. The typical durations for those activities are derived from the microcensus data and are specified individually for each member of the synthetic population. Based on further data, an initial mode choice was calculated with the restriction, that each agent can only use one transport mode for a plan.

The initial demand used for the simulations (referred to as "demand version 1" in the following) is based on the aforementioned demand of whole Switzerland, but consists only of all agents driving a car who, as part of their routing, are at least once inside an imaginary boundary around Zurich during their day. The boundary is defined as a circle with a radius of 30 kilometers ($\approx$ 18.6 miles) and with its center at "Bellevue", a central place in the city of Zurich. In order to obtain a higher computational speed, a random 10% sample was chosen for simulation, which consists of 61 480 agents.

The "default" strategy setup uses time adaptation and route adaptation. This means that in each iteration of the simulation, 10% of the agents adapt routes, while another 10% of the agents adapt activity times. The remaining 80% of the agents can select another plan among the plans in their plans collection. All modifications are reported with respect to this default strategy setup.

Activity type	Opening time	Closing time
Home	00:00	24:00
Work	06:00	20:00
Education	06:00	20:00
Shop	08:00	20:00
Leisure	00:00	24:00

TABLE 1 Activity opening and closing times used in the scenario.

Activity locations were given opening and closing times in order to keep the agents within some timely limit. The opening and closing times are classified by activity type, i.e. the opening and closing times are distinguished for home, work, education, shop and leisure activities. There is not yet any distinction based on the location of an activity. Table 1 summarizes the opening and closing times available to perform activities. Table 2 shows the behavioral parameters used in the scenario.

All simulations described in this paper are run for 500 iterations to retrieve a relaxed state, in which the initial plans are adapted to the traffic conditions. If not otherwise specified, time and route adaptation are enabled, each for 10% of the agents in every iteration.

VALIDATION METHODOLOGY

Modifications to network and travel demand only make sense if they help to increase the accuracy of the simulation results. This means that some kind of measurement must exist to determine the

Parameter	Value
β_{perf}	6 Euro /h
$\beta_{early.dp}$	-18 Euro /h
β_{car}	-6 Euro /h
$\beta_{non-car}$	-3 Euro /h
β (existing plans)	4

TABLE 2 Behavioral parameters used in the scenario.

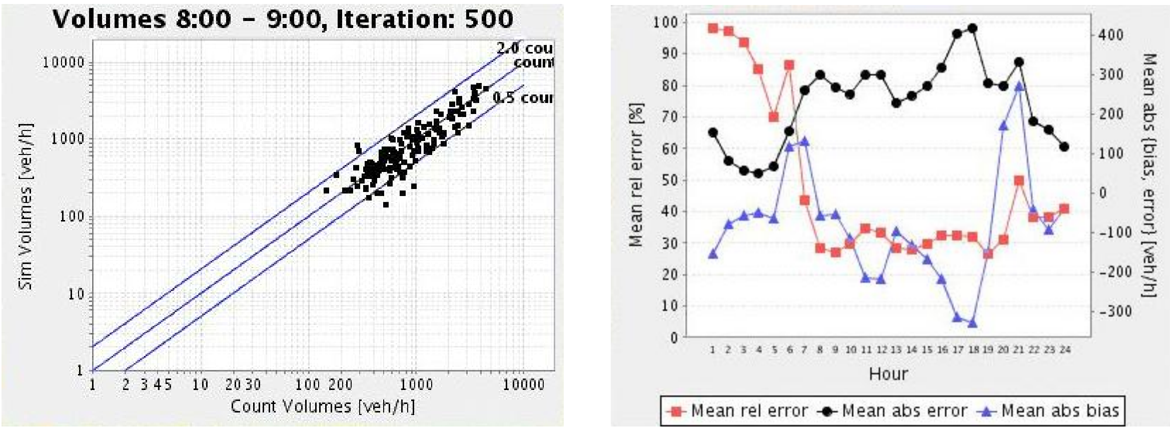

(a) Comparison for one hour

(b) Hourly analysis over time of day

FIGURE 1 Realism of the base case. 159 traffic counting stations provide real traffic counts for the Zurich area.

quality of the simulation. For the Zurich region, hourly data from 159 traffic counting stations is available. This data is used to compare the traffic volumes from the simulation to real-world values. Different statistical values can be calculated, like mean relative error or mean absolute bias. Fig. 1 shows two examples of standard reports that MATSim can automatically generate.

The model improvements described in this paper are all done to minimize the mean relative error (red curve in Fig. 1b). No formal decision was taken of how to weigh the different hours; instead, the graphs are interpreted visually.

The mean relative error for every sensor and every hour is calculated as:

$$\frac{\text{Simulated traffic volume} - \text{Real traffic volume}}{\text{Real traffic volume}} \ . \tag{1}$$

Averages for a given hour are obtained by averaging over all sensors. In the example shown in Fig. 1b, the simulation deviates strongly from the reality during the night hours, i.e. from midnight until 06:00 am. However, during daytime, i.e. from 06:00 am until late evening, the hourly mean relative error is around 30%.

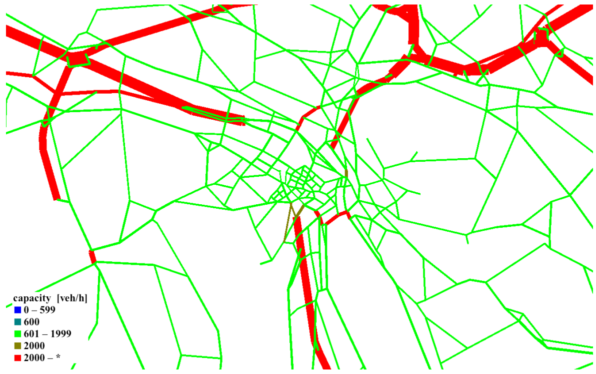

(a) original network (version 1)

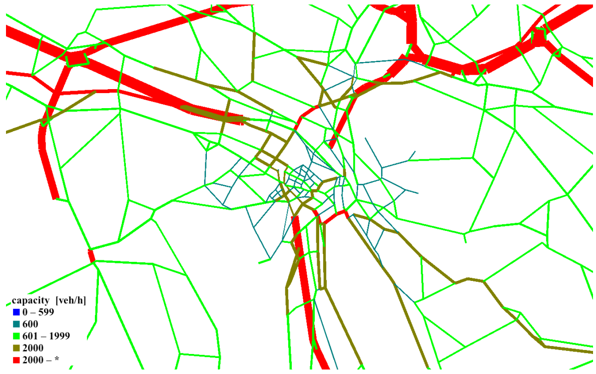

(b) modified network (version 2)

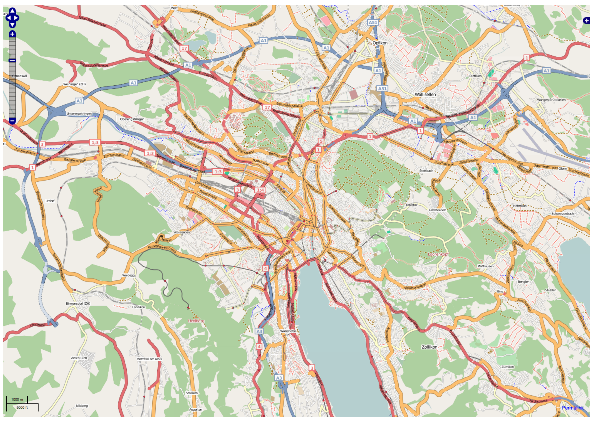

(c) network from www.openstreetmap.org

FIGURE 2 Link capacities in the original and modified network

IMPROVEMENTS

Two possible improvements are better network data and better travel demand. Both parts are essential for a realistic scenario. Additionally, it turns out that adding choice dimensions improves the quality of the simulation as well.

Network Improvements

The link capacities in the original network ("network version 1") are quite undifferentiated for most of the non-freeways (see Fig. 2a). The reason for this is most probably that the original network data is meant for Swiss-wide national analysis, and therefore a secondary network with a capacity that is approximately correct in the average is sufficient. Clearly, since we are interested in a better resolution at the urban scale, this is not sufficient.

To correct this problem, all links within a circle with radius 4 kilometers ($\approx$ 2.5 miles) around the center of Zurich are modified as follows:

- links corresponding to primary roads in OpenStreetMap (see Fig. 2c) get a capacity of at least 2 000 vehicles per hour. If the original capacity is higher than that, the capacity is not changed.
- links corresponding to secondary roads in OpenStreetMap keep their original capacity (usu-

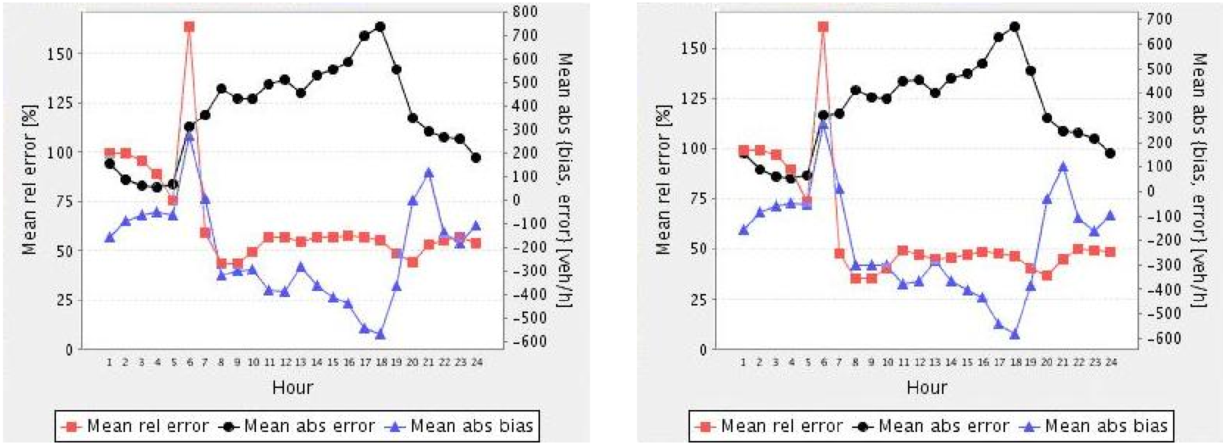

(a) original network (version 1) (b) modified network (version 2)

FIGURE 3 Error plots of simulations with the original and modified network

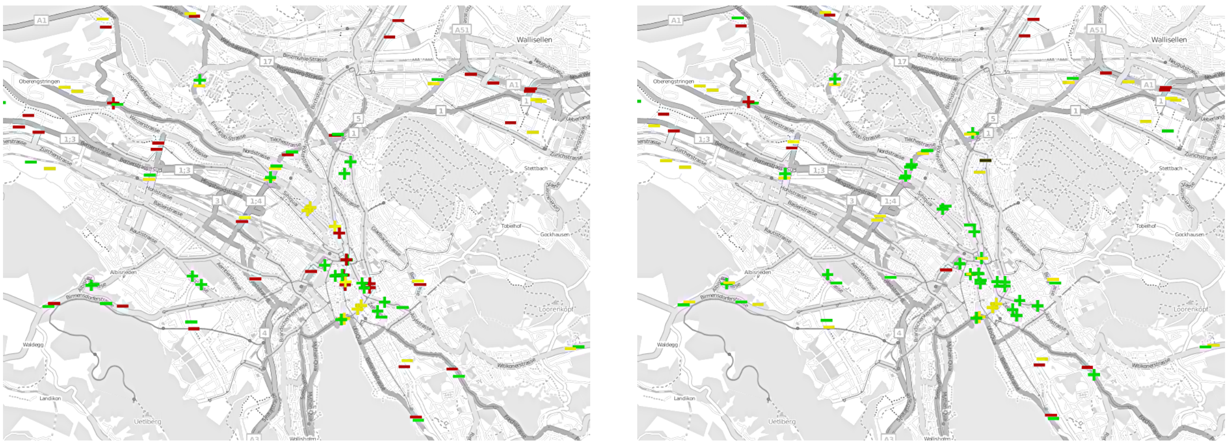

(a) original network (version 1) (b) modified network (version 2)

FIGURE 4 Locations of counting stations and comparison quality for the hour from 8am to 9am. Red symbols show a strong deviation of simulation volumes to counts, green symbols a good correlation. (Background-map: www.openstreetmap.org)

ally between 1 000 and 2 000 veh/h).

- all other links get a capacity of at most 600 veh/h. If the original capacity is lower, it is not changed.
- a few single links are manually adjusted based on local knowledge.

Fig. 2b shows the overview of the capacity of the links in this updated network ("version 2").

Comparing the mean relative errors from simulations with the two networks and the default strategy setup, one can see that the simulation with the adjusted network (version 2) has clearly a smaller mean relative error after 6am than the initial network (version 1) has (see Fig. 3a and Fig. 3b), with the mean relative error being now constantly below 50% during the daytime.

Fig. 4a and Fig. 4b show the geographical places where the counting stations are located on the network. The symbols and colors visualize the direct comparison of simulation volumes to the real-world volumes for one specific hour. This allows to relate under- or overestimated links to geographical characteristics. In Fig. 4a one can recognize that on most counting stations in the city center the traffic flows are overestimated by the simulation (symbolized with $+$), while outside the city center the traffic volumes on many links are underestimated (symbolized with $-$). Red symbols depict a strong deviation between simulated and real traffic volumes, while green symbols stand for no or only a small difference. Comparing the number of red symbols in both figures, one can see that their number is highly reduced in the network version 2, also proving the effectiveness of the network modifications. – All runs in this paper, with the exception of the run used for Fig. 2a, are done with the improved network ("network version 2").

As a side remark, runs with a network that is entirely based on physical characteristics plus time-of-day dependent green splits on intersections seems to perform even better (M. Balmer, personal communcation).

Demand Improvements

As described in the section "Scenario Setup", the original travel demand ("demand version 1") consists of agents traveling within the boundaries of Switzerland. When comparing the traffic volumes from the simulation with real-world data, one can observe that counting stations with too low volumes in the simulations are located especially along freeways, but only rarely in the city center or on smaller roads.

Further analysis resulted in the knowledge that a not to be underrated part of traffic on the freeways comes from abroad. Because of the short distance to neighboring countries (e.g. the border to Germany is less than 25 kilometers / 15 miles north of Zurich) it is not uncommon for people to live abroad but work in the Zurich area, or live in Zurich with its high cultural offers and work abroad. Those people are not part of the initial demand, as at least one of their activity locations lies outside Switzerland.

In addition, some of the intereuropean routes connecting Germany with Italy also pass through the greater area of Zurich. This leads to additional traffic not yet accomodated in the initial demand.

Both cases could be solved by adding "boarder-crossing traffic" (sometimes also referred as "through" traffic) *(24)*. Taking a 10% sample of all through traffic traveling with cars in the area of Zurich added 5759 agents to the demand. Running the simulation with this extended demand ("version 2") on network version 2 resulted in a clear improvement of the quality of the simulation, as Fig. 5 compared to Fig. 3b shows.

Improvements by adding choice dimensions

While network and demand are essential for realistic scenarios, the capabilities of the simulation itself also have a big influence on the quality of the results. To demonstrate that, the currently best case (network version 2, with modified capacities according to OpenStreetMap, and demand version 2, which includes through traffic) were run with different simulation features switched on or off:

- Route choice only, i.e. no mode choice, no time adaption. 10% of the agents can adapt their route in every iteration.
- Route choice *and* departure time choice, i.e. no mode choice.
- Route choice, departure time choice *and* mode choice.

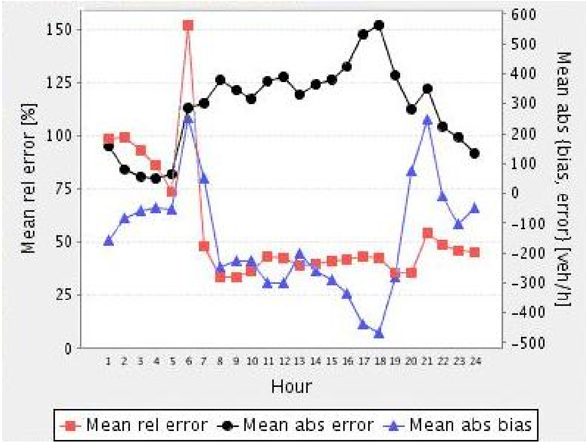

FIGURE 5 Error plots of the simulations with different travel demands

The first two cases, route-choice only and route- and time-choice, use the demand version 2, consisting of private car traffic within Switzerland and through traffic, as described in the section before. In these cases, the initial mode choice was used to determine which agents where driving a car and which ones not, and this remains fixed during the runs. In the third case where mode-choice is added, the initial mode choice is ignored. Instead, all agents from within Switzerland (= "demand version 1"; thus excluding those added by the through traffic) are given two plans, one where "car" is set as transport mode, and another one where "non-car" is set as transport mode. In all other aspects, the two plans are identical and identical to the plans in the original demand. This allows the agents to choose between the two transport modes, effectively adding mode choice to the scenario. Since now all agents are simulated and not only those with the initial mode choice set to "car", the number of simulated agents increases to 187 484.

Fig. 6 shows the quality of the different setups after 500 iterations. The more choice dimensions are available, the smaller is the mean relative error (red) for the time range from 8 am to 8 pm. For some other time segments, e.g. from 5 am to 6 am and from 8 pm to 9 pm, the mean relative error deviates from this tendency. This depends on the setting of activity opening and closing times (see Table 1); for a discussion of this see the "discussion" section.

It should be noted that the improvement from Fig. 6a to Fig. 6c occurs in spute of the fact that the initial demand is rather good, since the available data is rather detailed. The simulation is still able to improve the quality by determining its own time and mode choices. This could be explained with the fuzziness created by aggregations and extrapolations to apply the initial mode choice, while in the simulation the particular characteristics of the daily plan and the particular traffic characteristics of the car mode are automatically included.

It can be noted that in the case with route, time and mode choice, the mean relative error for the day hours is around 30%, sometimes a bit lower, sometimes a bit higher. While this may still be quite a big error in traditional, static traffic analysis applications, it is amongst the best results we know for dynamic traffic simulations, also because it is not only valid for the rush hours, but along a big part of the day.

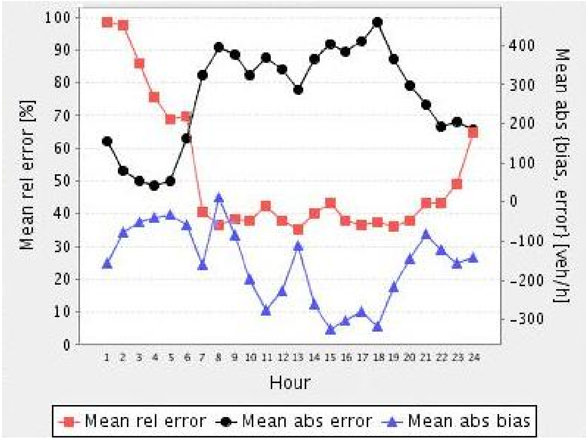

(a) route choice only

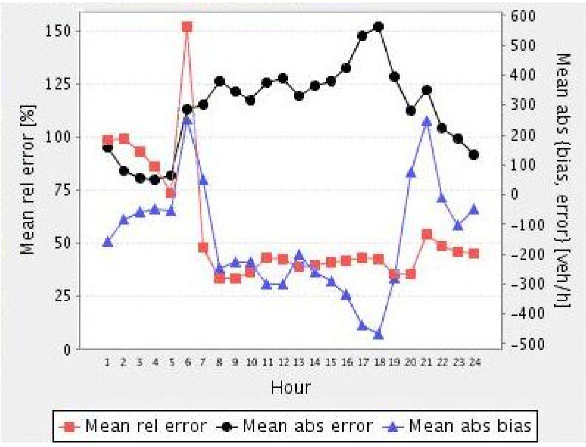

(b) route and time choice

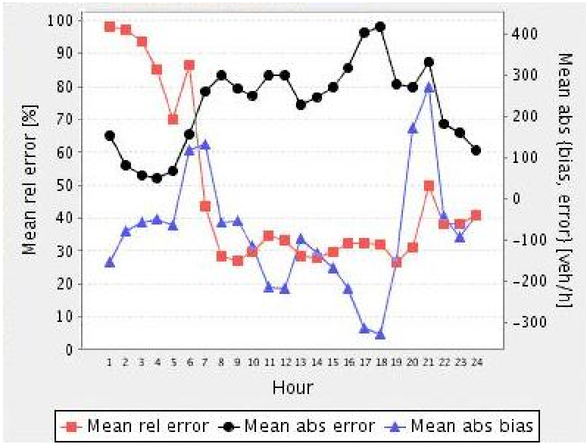

(c) route, time and mode choice

FIGURE 6 Error plots of the simulations with different replanning strategies

DISCUSSION

It was stated in the introduction that we are not aware of a currently existing practically feasible method to calibrate large scale models that use microscopic, behavior-based particles (= agents). Accordingly, the work presented above proceeds by heuristically adjusting different elements of the simulation. Still, we believe that the adjustments made in this paper can be defended: The improvement of the network was based on additional, and verifiably better input data; the improvement of the demand was based on demand that was clearly not covered by the existing demand generation; and it is in our view a good omen for the microsocpic method that opening up the choice dimensions made the model *better* and not worse.

We also believe that the facility opening times, encoded in Tab. 1, are defensible although the particular values might still be improved. In fact, the high peak of the mean relative error in most of the figures is due to the fact that the morning peak sets in too early on our simulations. This is also visible in the mean absolute bias, which shows a strong positive value in this time period. This is most probably due to work and education opening times which are too early (at 6am, see Tab. 1). Also, in reality not all facilities open at the same time. A similar, albeit weaker, effect seems to be at work at 8pm. Finding better values here should be done in future work.

An arguably more sensitive issue is the scoring function (utility function) which drives the adaptation of the agents. The assumption of the existence of a utility function in itself is already a strong statement *(25, 26)*, and this is important to state since the microscopic behavior-based approach could, in theory, also be based on alternative principles (see some of the chapters in *(27)*). If, however, we accept the utility function, it is instructive to look at its elements. As long as the synthetic travelers cannot drop activities, the following parameters are important:

- The typical duration, $t_{*,i}$ of each activity type.
- The slope of the utility function at its typical duration, for each activity type.
- The curvature of the utility function at its typical duration, for each activity type.

In consequence, at first glance it seems that there are three free parameters per activity type. Fortunately, this number can be reduced by the following arguments:

- In order to be optimal, the activity durations need to be selected such that all slopes (= marginal utilities) are the same, at least in the absence of constraints such as opening times or other influences such as strongly variable travel times. This implies that one can, as a first approximation, set all slopes at the typical duration to the same value. This ends up being the marginal utility of leisure time *(28)*, which can be estimated.
- By the same argument, it should be possible to estimate "typical durations" of activity times from time use surveys: If marginal utilities are the same, then the typical durations need to be set such that the typical durations from time use surveys are recovered. – In our current work, the typical durations are directly taken from actual durations from time use surveys in Switzerland (see below).

The curvature at the typical durations remains as the most problematic parameter. This parameter determines the flexibility of an activity: a large curvature means that the marginal utility increases strongly when the activity duration is reduced, implying that time should rather be saved somewhere else. For the time being, however, our choice of a function of type $A \ln(t/t_0)$ has the consequence that, after the slope at the typical duration has been set to the correct value as indicated above, no additional free parameter is left and, in consequence, there is nothing to calibrate.[2]

CONCLUSION

A microscopic, behavior-based ("agent-based") traffic simulation model is applied to the region of Zurich in Switzerland. The model is validated against hourly traffic counts of 159 counting stations. Not surprisingly, better demand data and better network data leads to better results. What is a bit more surprising is that the network improvements were informed by the freely available openstreetmap data, nourishing some hope that important network data for traffic modelling may eventually become available through this internationally available and standardized data source.

However, the most important result is that adding choice dimensions to the simulation (from "route choice only" to "route and time choice" and finally to "route, time, and mode choice") makes the results more realistic. Our interpretation is that the locally optimizing agents are able to pick up local particularities of the urban system that are missed by more aggregate methods.

[2]The results presented in this paper use, in addition, a minimum duration penalty, i.e. if an activity is left before its typical duration is up, there is an additional penalty. This has a similar effect as a strong curvature at the typical duration, but it is rather ad-hoc, and we will investigate the effect of this penalty.

This is, in our view, a good omen for the microscopic, behavior-based methods which have always made the claim that, in the end, they might be more parsimonious than other methods.

ACKNOWLEDGMENTS
We would like to thank Michael Balmer and the group of Prof. Axhausen at the IVT, ETH Zurich, for their work that led to the initial network and demand as described in the section “Scenario setup”. Furthermore we greatly acknowledge the work done by the uncounted contributors to the project OpenStreetMap for their data acquisition and data processing. The simulations were run on our compute cluster, set up and maintained by the Department of Mathematics at the Berlin Institute of Technology (TU Berlin), administered by the group of Prof. Schwandt.

REFERENCES

1. J. Jonnalagadda, N. Freedman, W.A. Davidson, and J.D. Hunt. Development of microsimulation activity-based model for San Francisco: destination and mode choice models. *Transportation Research Record*, 1777:25–35, 2001.

2. J.L. Bowman, M. Bradley, Y. Shiftan, T.K. Lawton, and M. Ben-Akiva. Demonstration of an activity-based model for Portland. In *World Transport Research: Selected Proceedings of the 8th World Conference on Transport Research 1998*, volume 3, pages 171–184. Elsevier, Oxford, 1999.

3. P. Vovsha, E. Petersen, and R. Donnelly. Microsimulation in travel demand modeling: lessons learned from the New York best practice model. *Transportation Research Record*, 1805:68–77, 2002.

4. D. Boyce and H. Bar-Gera. Validation of multiclass urban travel forecasting models combining origin-destination, mode, and route choices. *Journal of Regional Science*, 43:517–540, 2003.

5. O.A. Nielsen and R.D. Frederiksen. Optimisation of timetable-based, stochastic transit assignment models based on MSA. *Annals of Operations Research*, 144:263–285, 2006.

6. D. Lohse, C. Schiller, H. Teichert, M. Vrtic, P. Fröhlich, N. Schüssler, and K.W. Axhausen. Ein zweiseitig gekoppeltes Modell zur simultanen Berechnung der Verkehrserzeugung, Verkehrsverteilung und Verkehrsaufteilung: theoretischer Hintergrund und praktische Anwendung für ein nationales Modell der Schweiz. *Verkehrsforschung Online*, 3(1):1, 2006.

7. M. Balmer, K.W. Axhausen, and K. Nagel. A demand generation framework for large scale micro simulations. *Transportation Research Record*, 1985:125–134, 2006.

8. K. Meister, M. Rieser, F. Ciari, A. Horni, M. Balmer, and K.W. Axhausen. Anwendung eines agentenbasierten Modells der Verkehrsnachfrage auf die Schweiz. In *Proceedings of Heureka '08*, Stuttgart, Germany, March 2008.

9. BIOGEME www page, accessed 2008.

10. G. Flötteröd. *Traffic state estimation with multi-agent simulations*. PhD thesis, Berlin Institute of Technology, 2008.

11. G. Flötteröd and K. Nagel. State estimation for multi-agent simulations of traffic. In *Proceedings of the Word Conference on Transport Research*, Berkeley, CA, 2007.

12. B. Raney and K. Nagel. An improved framework for large-scale multi-agent simulations of travel behaviour. In P. Rietveld, B. Jourquin, and K. Westin, editors, *Towards better performing European Transportation Systems*, page 42. Routledge, London, 2006.

13. MATSIM www page. MultiAgent Transport SIMulation. http://matsim.org/, accessed 2008.

14. K. Nagel, M. Strauss, and M. Shubik. The importance of timescales: Simple models for economic markets. *Physica A*, 340(4):668–677, 2004.

15. N. Lefebvre and M. Balmer. Fast shortest path computation in time-dependent traffic networks. In *Proceedings of Swiss Transport Research Conference (STRC)*, Monte Verita, CH, September 2007. See www.strc.ch.

16. C. Gawron. *Simulation-based traffic assignment*. PhD thesis, University of Cologne, Cologne, Germany, 1998.

17. N. Cetin, A. Burri, and K. Nagel. A large-scale agent-based traffic microsimulation based on queue model. In *Proceedings of Swiss Transport Research Conference (STRC)*, Monte Verita, CH, 2003. See www.strc.ch. Earlier version, with inferior performance values: Transportation Research Board Annual Meeting 2003 paper number 03-4272.

18. D.E. Kaufman, K.E. Wunderlich, and R.L. Smith. An iterative routing/assignment method for anticipatory real-time route guidance. Technical Report IVHS Technical Report 91-02, University of Michigan Department of Industrial and Operations Engineering, Ann Arbor MI 48109, May 1991.

19. J.A. Bottom. *Consistent anticipatory route guidance*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 2000.

20. E. Avineri and J.N. Prashker. Sensitivity to uncertainty: Need for paradigm shift. *Transportation Research Record*, 1854:90–98, 2003.

21. D. Charypar and K. Nagel. Generating complete all-day activity plans with genetic algorithms. *Transportation*, 32(4):369–397, 2005.

22. M. Vrtic, P. Fröhlich, and K.W. Axhausen. Schweizerische Netzmodelle für Strassen- und Schienenverkehr. In T. Bieger, C. Lässer, and Rico Maggi, editors, *Jahrbuch 2002/03 Schweizerische Verkehrswirtschaft*, pages 119–140. Schweizer Verkehrswissenschaftliche Gesellschaft, St. Gallen, 2003.

23. SFSO. Ergebnisse des Mikrozensus 2005 zum Verkehr. Swiss Federal Statistical Office, Neuchatel, 2006.

24. M. Vrtic, N. Schüssler, A. Erath, K. Meister, and K. Axhausen. Tageszeitliche Fahrtenmatrizen im Personenverkehr an Werktagen im jahr 2000. Research report, Swiss Federal Department for Environment, Transport, Energy and Communication, Swiss Federal Office for Spatial

Development, Swiss Federal Roads Authority and Swiss Federal Office for Transport, IVT, ETH Zürich, Zürich, 2007.

25. H.A. Simon. *An empirically based microeconomics*. Cambridge University Press, 1997.

26. S. Moss and E.-M. Sent. Boundedly versus procedurally rational expectations. In A.H. Hallett and P. McAdam, editors, *Analyses in macroeconomic modelling*, Advances in computational economics, pages 115–146. Springer, Heidelberg, 1999.

27. Th. Arentze and H.J.P. Timmermans, editors. *ALBATROSS–Version 2.0 – A learning based transportation oriented simulation system*. EIRASS (European Institute of Retailing and Services Studies), TU Eindhoven, NL, 2005.

28. S. Jara-Díaz, M.A. Munizaga, P. Greeven, and P. Romero. The activities time assignment model system: The value of work and leisure for Germans and Chileans. In *Presented at the European Transport Conference (ETC)*, 2004.