RESEARCHING THE INFLUENCE OF TIME-DEPENDENT TOLLS WITH A MULTI-AGENT TRAFFIC SIMULATION

Marcel Rieser
Transport Systems Planning and Transport Telematics, TU Berlin, and
Institute for Transport Planning and Systems, ETH Zurich
Ulrike Beuck, Kai Nagel
Transport Systems Planning and Transport Telematics, TU Berlin

INTRODUCTION

In many cities in Europe, traffic tolls are discussed as a means to reduce the amount of traffic during peak hours or in general in a city. But only very few cities are able to indeed test at least one toll scheme, as the introduction of a toll is coupled to high investments and is not likely to be a popular activity for the politicians. Thus many cities have a big interest in transportation planning tools and traffic models to thoroughly test different toll schemes in a model to find one that solves their problem best—including acceptance in the population.

Traditional transportation planning tools work macroscopically, distributing a static traffic flows onto a network. While this is a well-established technology, it is not able to fully model all aspects that are of interest when modelling tolls. In particular, they usually lack any meaning of time-of-day. The models usually calculate the traffic flows for a complete day, or at best for certain periods (morning peak, evening peak), but in all cases without any temporal development. This makes it difficult to model time-dependent tolls, as the reaction of the travellers (e.g. driving before/after the toll) cannot be modelled by the planning tools, but must be given by the user. This reduces the usefulness of such a tool enormously.

Dynamic traffic assignment (DTA) explicitly models the temporal development of the traffic. Demand, however, is typically given as fixed-period (e.g. hourly) OD matrices, and does, in consequence, not adapt to the toll. Adaptation would need to happen in the demand generation modules that *generate* the OD matrices, but that implies rather intricate coupling between demand generation and DTA. In addition, the DTA is no longer aware of traveller characteristics, such as income or time constraints, and can*not*, in consequence, base any kind of toll route acceptance/rejection decision on such attributes.

A partial way it is METROPOLIS (de Palma and Marchal 2002), which selects departure times of trips based on desired arrival times and schedule delay penalties. Given a time-dependent toll, travellers can react by selecting different departure times. A remaining problem is, however, the fact that trips and in consequence decisions are not related to demographics. In addition, every model that uses single trips only will have problems predicting useful reactions of travellers that span the whole day. This is because trips in real life are embedded in a complete day plan and are not meaningful just as stand-alone trips. Trips lead people from one activity to another, and in most cases the activities have a higher importance in the daily schedule than the trips do: Stores we want to go to for shopping have opening and closing times, work places have fixed times where one has to be present, a regular worker has to work about eight hours a day. This means that travellers cannot escape a toll at their will, but have to trade off between different utilities (working eight hours, being at a shop when it has opened, …) and disutilities (paying a toll, being late for work, …). Thus a toll may influence the whole day schedule of a person, and not only the duration the toll is active.

Our approach uses multi-agent simulations to model and simulate full daily plans. This allows us to research the influence of time-dependent tolls more thoroughly than traditional tools are able to.

SIMULATION STRUCTURE

Overview

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 (and sometimes the 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. Further details will be given below.

The simulation approach is the same as in many of our previous papers (e.g. Raney and Nagel 2006, Balmer et al 2005) on the same subject. The following exposition thus borrows heavily from those papers. In addition, the results of this paper are based on a re-implementation of the MATSim framework in Java (MATSim 2007). This has made the computational performance of the code somewhat slower (in particular, distributed computing is no longer supported), but allows faster conceptual progress.

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. This paper concentrates on "home" and "work" as the only activities, and "car" as the only mode. Our implementation already at this point supports additional activity types (see, e.g., Meister et al 2006) and additional modes of transport, but more time is needed to validate results with those additional complexities.

The task of generating a plan is divided into sets of decisions, and each set is assigned to a separate module. An agent strings together calls to various modules in order to build up a complete plan. To support this “stringing”, the input to a given module is a (possibly incomplete) plan, and the output is a plan with some of the decisions updated. This paper will make use of two modules only: “activity times generator” and “router”. Other modules will be the topic of future work. Once the agent’s plan has been constructed, it can be fed into the traffic flow simulation. This module executes all agents’ plans simultaneously on the network, allowing agents to interact with one another, and provides output describing what happened to the agents during the execution of their plans.

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). This creates an interdependency (“chicken and egg”) problem between the decision-making modules and the traffic flow simulation. To solve this, feedback is introduced into the multi-agent simulation structure (Kaufman et al 1991, Bottom 2000). 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; 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.

In the following sections we describe the used modules in more detail.

Activity Time Allocation Module

This module is called to change the timing of an agent’s plan. At this point, a very simple approach is used which applies a random “mutation” to the duration and end time attributes of the agent’s activities. For each such attribute of each activity in an agent’s plan, this module picks a random time from the uniform distribution [–30 min, +30 min] and adds it to the attribute. Any negative duration is reset to zero; any activity end time after midnight is reset to midnight.

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. Nagel et al 2004). Since each module is implemented as a “plugin”, this module can be replaced by a more enhanced implementation if desired.

MATSim contains already a more sophisticated activity scheduling module (Meister et al 2006). This will be used in future studies.

Router

The router is implemented as a time dependent Dijkstra algorithm. It calculates link travel times from the events output of the previous traffic flow simulation (see next section). 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. Apart from relatively small and essential technical details, the implementation of such an algorithm is straightforward (Jacob et al 1999). With this and the knowledge about activity chains, it computes the fastest path from each activity to the next one in the sequence as a function in time.

Traffic Flow Simulation

The traffic flow simulation simulates the physical world. It 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 (Gawron 1998, Cetin et al 2003). 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.

Even though this structure is indeed very simple, it produces traffic as expected and it can run directly off the data typically available for transportation planning purposes. On the other hand, there are some limitations compared to reality, i.e. number of lanes, weaving lanes, turn connectivities across intersections or signal schedules cannot be included into this model.

The output that the traffic flow simulation produces is a list of events for each agent, such as entering/leaving link, left/arrived at activity, and so on. Data for an event includes which agent experienced it, what happened, what time it happened, and where (link/node) the event occurred. With this data it is easy to produce different kinds of information and indicators like link travel time (which i.e. will be used by the router), trip travel time, trip length, percentage of congestion, and so on.

Agent Database – Feedback

As mentioned above, the feedback mechanism is important for making the modules consistent with one another, and for enabling agents to learn how to improve their plans. In order to achieve this improvement, agents need to be able to try out different plans and to tell when one plan is “better” than another. The iteration cycle of the feedback mechanism allows agents to try out multiple plans. To compare plans, the agents assigns each plan a “score” based on how it performed in the traffic flow simulation.

It is very important to note that our framework always uses actual plans performance for the score. This is in stark contrast to all other similar approaches that we are aware of. These other approaches always feed back some aggregated quantity such as link travel times and reconstruct performance based on those (e.g. URBANSIM 2007, Ettema et al 2003). Because of unavoidable aggregation errors, such an approach can fail rather badly, in the sense that the performance information derived from the aggregated information may be rather different from the performance that the agent in fact displayed (Raney and Nagel 2004).

The procedure of the feedback and learning mechanism is described in detail in (Balmer et al 2005). For better understanding, the key points are restated here.

The agent database starts with one complete plan per agent, which is marked as "selected". The simulation executes these marked plans simultaneously and outputs events. Each agent uses the events to calculate the score of its "selected" plan and decides, which plan to select for execution by the next traffic flow simulation. When choosing a plan, the agent database can either:

• create a new plan by sending an existing plan to the router, adding the modified plan as a new plan and selecting it,

• create a new plan by sending an existing plan to the time allocation module, adding the modified plan and selecting it,

• pick an existing plan from memory, choosing according to probabilities based on the scores of the plans. The probabilities are of the form $p \propto e^{\beta \cdot S_j}$, where S_j is the score of plan j, and β is an empirical constant. This is similar to a logit model from discrete choice theory.

After this step, the simulation executes the newly selected plans again. 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.

Scores for plans

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, Avineri and Prashker 2003). In this work, a simple utility-based approach is used. The approach is related to the Vickrey bottleneck model (Arnott et al 1993), but is modified in order to be consistent with our approach based on complete daily plans (Charypar and Nagel 2005, Raney and Nagel 2006). The elements of our approach are as follows:

• the total score of a plan is computed as the sum of individual contributions:

$$U_{total} = \sum_{i=1}^{n} U_{perf,i} + \sum_{i=1}^{n} U_{late,i} + \sum_{i=1}^{n} U_{travel,i}$$

where U_{total} is the total utility for a given plan; n is the number of activities, which equals the number of trips; $U_{perf,i}$ is the (positive) utility earned for performing activity i; $U_{late,i}$ is the (negative) utility earned for arriving late to activity i; and $U_{travel,i}$ is the (negative) utility earned for traveling during trip i. In order to work in plausible real-world units, utilities are measured in Euro.

• 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.

$t_{0,i}$ is a scaling parameter that is related both to the minimum duration and to the importance of an activity. If the actual duration falls below $t_{0,i}$, then the utility contribution of the activity becomes negative, implying that the agent should rather completely drop that activity. A $t_{0,i}$ only slightly less than $t_{*,i}$ means that the marginal utility of activity i rapidly increases with decreasing $t_{perf,i}$, implying that the agent should rather cut short other activities. This paper uses $t_{0,i} = t_{*,i} \cdot \exp(-\xi / t_{*,i})$ where ξ is a scaling constant set to 10 hours. With this specific form, $U_{perf,i}(t_{*,i}) = \beta_{perf} \cdot \xi$ is independent of the activity type.

• The (dis)utility of being late is uniformly assumed as: $U_{late,i} = \beta_{late} \cdot t_{late,i}$, where β_{late} is the marginal utility (in Euro/h; usually negative) for being late, and $t_{late,i}$ is the number of hours late to activity i.

• The (dis)utility of traveling is uniformly assumed as: $U_{travel,i} = \beta_{travel} \cdot t_{travel,i}$, where β_{travel} is the marginal utility (in Euro/h; usually negative) for travel, and $t_{travel,i}$ is the number of hours spent traveling during trip i.

In principle, arriving early or leaving early could also be punished. There is, however, no immediate need to punish early arrival, since waiting times are already indirectly punished by foregoing the reward that could be accumulated by doing an activity instead (opportunity cost). In consequence, the effective (dis)utility of waiting is already $-\beta_{perf}$. Similarly, that opportunity cost has to be added to the time spent traveling, arriving at an effective (dis)utility of traveling of $-|\beta_{travel}| - \beta_{perf}$.

No opportunity cost needs to be added to late arrivals, because the late arrival time is spent somewhere else. In consequence, the effective (dis)utility of arriving late remains at β_{late} . – These values (β_{perf} , $\beta_{perf} + |\beta_{travel}|$, and $|\beta_{late}|$) are the values that would correspond to the consensus values of the parameters of the Vickrey model (Arnott et al 1993) if MATSim would just look for late arrival.

SCENARIO

The chosen scenario is the same as in (Balmer et al 2005). It covers the area of Zurich, Switzerland, which has about 1m inhabitants. The network is a Swiss regional planning network, extended with the major European transit corridors (figure 1.a). It has the fairly typical size of 10 564 nodes and 28 624 links.

The simulated demand consists of commuters only that travel by car in the aforementioned region, resulting in 260 275 agents, all with an activity pattern home-work-home. The initial time structure has the agents leaving home in the morning at a randomly chosen time between 6am and 9am, work for 8 hours, and then returning to home.

TOLLS

During the rush hours, traffic in Zurich is very dense, and tolls are debated as a possible solution. Thus, we defined a hypothetical toll area that covers the inner city of Zurich, but not the highways that lead into and partially around the city. Figure 1.b shows the area with the tolled links. The diameter of the toll area ranges from 5–7km. The toll is restricted to the evening (3pm to 7pm) only and is set to 2 Euro/km. This may sound steep, but we wanted a clear signal in this somewhat synthetic scenario. Restriction of the toll to the evening is done to illustrate that the agent-based approach as is able to consider ramifications throughout the whole day. In particular, it will be shown that the morning traffic is significantly affected by the evening toll. As was discussed earlier, this is an effect that a trip-based model cannot represent. The covered area has a high density of offices and other work places, so the in-bound traffic is larger in the morning than the out-bound traffic, and vice versa in the evening.

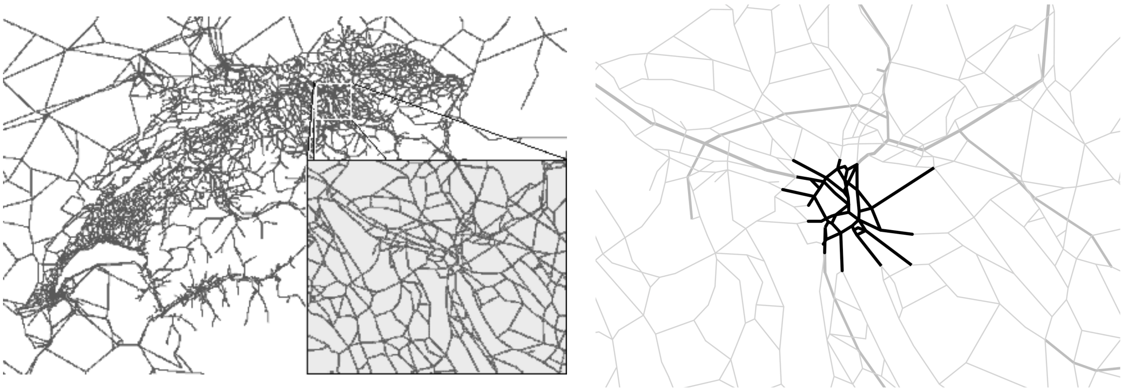

(a) Switzerland network, area of Zurich enlarged
(b) hypothetical toll links in Zurich area.
Figure 1: Scenario: Switzerland network with toll links for Zurich

RUNS

A base case without the toll was first iterated until a relaxed state was reached. Both routes and activity times were allowed to adapt. Based on this state, a new iterations run was started with the toll switched on, again until a (new) relaxed state was reached. This allows researching the specific influence the toll has on the behaviour of the travellers.

We used the following values for the marginal utilities of the utility function used for calculating scores:

β_{perf} = +6 Euro/h, β_{travel} = –6 Euro/h, β_{late} = –18 Euro/h

Although it is not obvious at a first glance, these values mirror the standard values of the Vickrey scenario (e.g. Arnott, de Palma, and Lindsey 1993): An agent that arrives early to an activity must wait for the activity to start, therefore forgoing the β_{perf} = +6 Euro/h that it could accumulate instead. An agent that travels forgoes the same amount, plus a loss of 6 Euro/h for traveling. And finally, an agent that arrives late receives a penalty of 18 Euro/h late.

In addition, this paper will only look at daily activity chains that consist of on home and one work activity. The optimal times will be set to
$t_{*,h}$ = 15 hours and $t_{*,w}$ = 8 hours

For the work activity a starting time window is defined between 7:08am and 8:52am. These values were set to correspond with those used in a similar study (17).

RESULTS

A first visual validation is done by looking at the traffic volumes and velocities. Figure 2 shows the velocity of agents at 5.30pm, during the toll hours. One can clearly see that there are more agents travelling outside the toll area in the toll case by the traffic jams they produce.

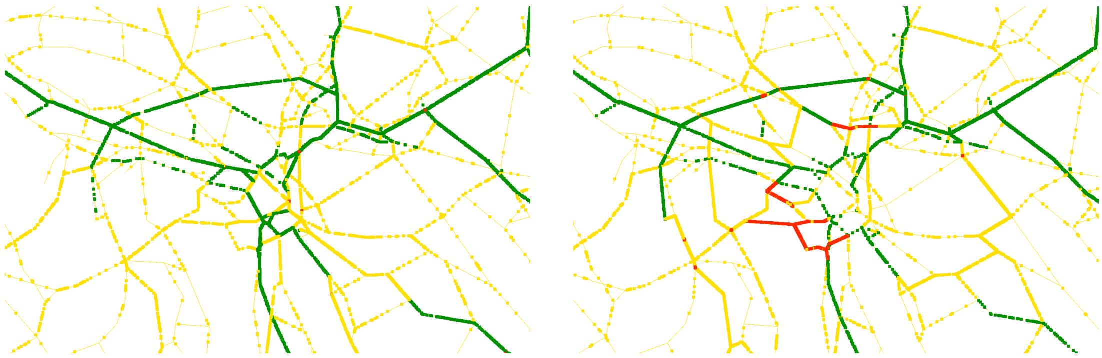

(a) base case (b) toll case

Figure 2: Travel speeds at 5.30pm during the toll time on the network.
Green are high speeds, red marks traffic jams.

We can also compare the two runs in the morning hours at 8am. Note that at this time of the day, there is no active toll in both cases! As can be seen in figure 3, there are traffic jams in the base case, but none of them in the toll case. This clearly shows that the toll in the evening rush hour has an influence on the morning rush hour.

Figure 4.a shows the departure time distribution at different times of day for the base case and the toll case. Comparing the toll case with the base case in the evening peak, one can see nicely how the number of travellers departing from work is higher in the toll case than in the base case in the time before the toll starts. It is also slightly higher after the toll ends. However, during the time the toll is active, the number of travellers departing is lower in the toll case than in the base case.

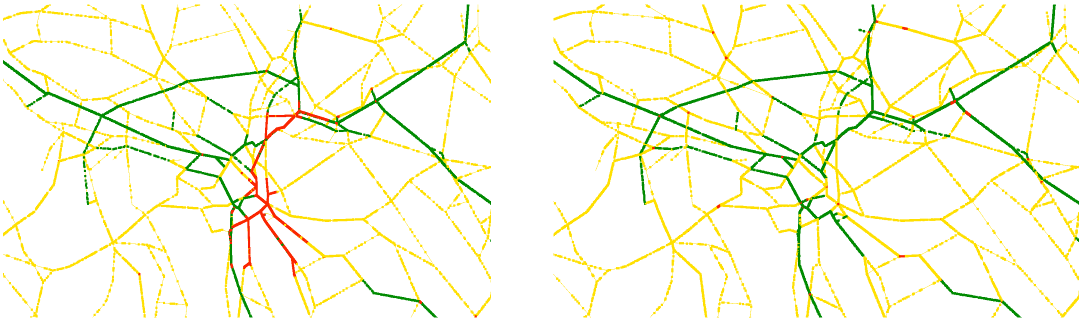

(a) base case (b) toll case

Figure 3: Travel speeds at 8am, when no toll has to be paid.
Green are high speeds, red marks traffic jams.

As each traveller tries to work eight hours a day, the same characteristics can also be seen in the morning rush hour, as agents planning to leave before 3pm will also have to arrive at work earlier than the others. This leads to a general broadening of the two peaks in the morning and the evening.

If the peaks of departing travellers are broader but less high, this means also that there are likely fewer people travelling at the same time. Figure 4.b shows the number of travellers simultaneously on the road. Especially in the morning rush hour it is apparent that the area below the curve is significantly smaller than in the base case. The area below the curve can be interpreted as the total time agents spend on the road. A smaller area means that people spend less time in total travelling—and all this without a toll in the morning rush hour!

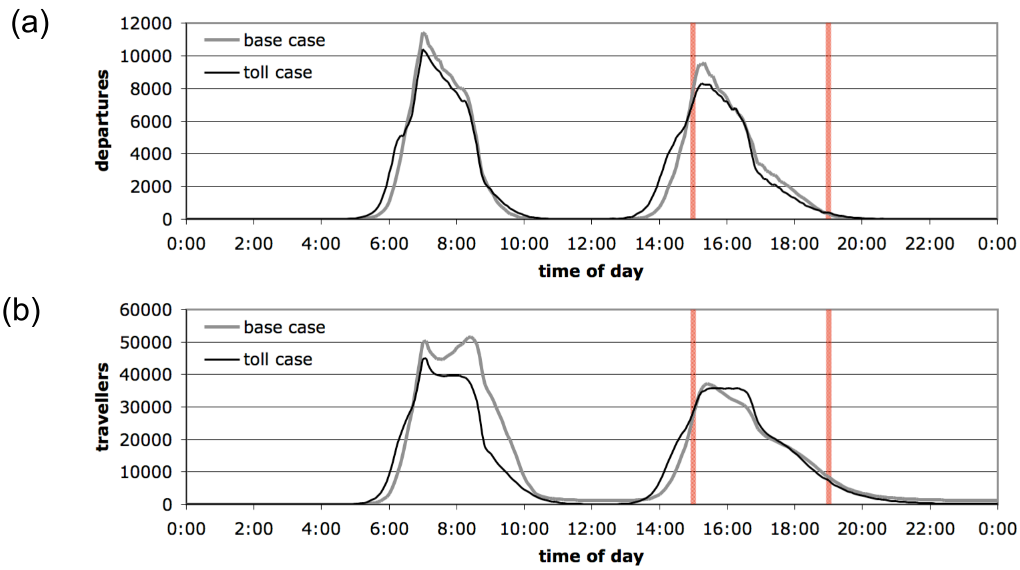

Figure 4: Number of departures (a) and number of travellers on the road (b) over the time of day.
The red lines mark the start and the end of the time a toll has to be paid.

The case is a bit different in the evening rush hour. Around 4pm we actually have more travellers on the road than in the toll case. This can be explained if one remembers that the toll area is only a small part of the whole simulated area: The travellers only have to get out of the toll area before the toll starts (as can be seen in the higher number of departures and travellers between

2pm and 3pm). However, this has the consequence that there may now be more travellers outside the toll area—and that's what can be observed in figure 4.b at 4pm.

One can also look at the number of travellers on a single link. For example, we can look at the number of travellers on an out-bound link (figure 5). One can see how the number of travellers on the link reaches a maximum at the beginning of the toll, where many agents try to cross the link before the toll starts. After that, the number of travellers decreases.

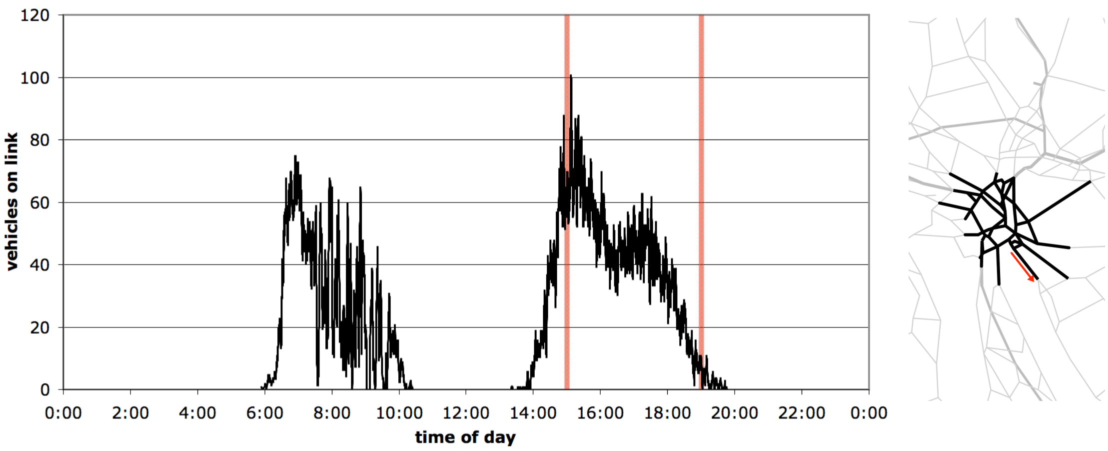

Figure 5: The number of vehicles on a single link. The red lines in the graph mark the start and the end of the time a toll has to be paid. The arrow in the small network is parallel to the selected link and shows the direction of the traffic on the link.

CONCLUSION

We have shown that multi-agent simulations can be used to model travellers' reactions to time-dependent toll in a way most existing transportation planning tools are not able to. As time-dependent tolls are a much-debated object in transportation politics, the ability to fully model such tolls and the reactions of travellers may help to find better toll schemes or to base the decision for or against a specific toll scheme more thoroughly. Additionally, the chosen multi-agent simulations are able to answer a multitude of additional questions current transportation tools are not able to. In a world where individuals have more and more freedom to schedule their daily plans, agent-based simulations offer an intuitive way to research complex topics with lots of interdependencies—like the interdependence of different trips for a single agent throughout the day.

BIBLIOGRAPHY

Arnott, R., De Palma, A., and Lindsey, R. (1993) A structural model of peak-period congestion: A traffic bottleneck with elastic demand, **The American Economic Review, 83** (1):161.

Avineri, E. and Prashker, J. N. (2003) Sensitivity to uncertainty: Need for paradigm shift. **Transportation Research Board Annual Meeting**, Washington, D.C.

Balmer, M., Raney, B., and Nagel, K. (2005) Adjustment of activity timing and duration in an agent-based traffic flow simulation. In H.J.P. Timmermans, editor, **Progress in activity-based analysis**, 91–114.

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

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

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

De Palma, A., and Marchal, F. (2002) Real case application of the fully dynamic METROPOLIS tool-box: An advocacy for large-scale mesoscopic transportation systems, **Networks and Spatial Economics**, **2** (4), 347–369.

Ettema, D., Tamminga, G., Timmermans, H., and Arentze, T. (2003) A micro-simulation model system of departure time and route choice under travel time uncertainty. **Proceedings** *of the meeting of the International Association for Travel Behavior Research (IATBR)*, Lucerne, Switzerland.

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

Kaufman, D. E., Wunderlich, K. E., and Smith, R. L. (1991) *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.

Jacob, R. R., Marathe, M. V., and Nagel, K. (1999) A computational study of routing algorithms for realistic transportation networks. **ACM Journal of Experimental Algorithms**, **4** (6).

MATSIM www page. MultiAgent Transport SIMulation, accessed 2007. URL www.matsim.org.

Meister, K., Balmer, M., Axhausen, K. W., and Nagel, K. (2006) planomat: A comprehensive scheduler for a large-scale multi-agent transportation simulation. **Proceedings** *of the meeting of the International Association for Travel Behavior Research (IATBR)*, Kyoto, Japan.

Nagel, K., Strauss, M., and Shubik, M. (2004) The importance of timescales: Simple models for economic markets. **Physica A**, 340:668–677, ISSN 0378-4371.

Raney, B. and Nagel, K. (2004) Iterative route planning for large-scale modular transportation simulations. **Future Generation Computer Systems**, **20** (7):1101–1118, ISSN 0167-739X.

Raney, B. and Nagel, K. (2006) An improved framework for large-scale multi-agent simulations of travel behavior. In P. Rietveld, B. Jourquin, and K. Westin, editors, **Towards better performing European Transportation Systems**. Routledge, Oxon, UK, 2006. Similar version TRB preprint number 05-1846.

URBANSIM www page, accessed 2007. URL www.urbansim.org.