MULTI-AGENT TRANSPORT SIMULATIONS AND ECONOMIC EVALUATION

Marcel Rieser
Transport Systems Planning and Transport Telematics, TU Berlin, and
Institute for Transport Planning and Systems, ETH Zurich

Ulrike Beuck, Dominik Grether, Yu Chen, Kai Nagel
Transport Systems Planning and Transport Telematics, TU Berlin

Kay W. Axhausen
Institute for Transport Planning and Systems, ETH Zurich

1. INTRODUCTION

In many cities in Europe, tolls are discussed as a means to reduce the amount of traffic during peak hours or in a city in general. But only few cities (e.g. Singapore, Bergen, Trondheim, Durham, London, Stockholm, Bologna, Roma) have been able to implement a toll scheme, as tolling requires substantial upfront investment and is likely to be an unpopular policy. Thus, many cities have an interest in transportation planning tools and traffic models to thoroughly test different toll schemes to find one that solves their problems best—and provides hints about popular acceptance.

Traditional transportation planning tools work macroscopically, distributing static traffic flows onto a network (e.g. Ortùzar and Willumsen 1995). While this is a well-established technology, it is not able to model all aspects that are of interest when modelling tolls. In particular, they usually lack any meaningful representation of the time-of-day dynamics. The models usually calculate traffic flows for a complete day, or at best for certain periods (morning peak, evening peak), but in all cases without any feedbacks between the different times-of-day. This makes it difficult or impossible to model time-dependent tolls, as the reaction of the travellers (e.g. driving before/after the toll) cannot be captured endogenously within the planning model, but must be pre-specified by the analyst. This reduces the usefulness of such a tool enormously.

Dynamic traffic assignment (DTA) explicitly models the temporal dynamics of traffic over the day (Chang et al 1985, Ben-Akiva et al 1998, Bliemer 2003). Demand, however, is typically given as fixed-period (e.g. hourly) origin-destination (OD) matrices, and, in consequence, does 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, in consequence, can*not* base any kind of toll route acceptance/rejection decision on such attributes.

A partial way out is offered by the approach chosen within 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 new 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 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 have opening and closing times, work places have fixed times when one has to be present, a full-time employee has to work about eight hours a day. This means that travellers cannot escape a toll at will, but have to trade off between different utilities (working eight hours, being at a shop when it has opened, etc.) and disutilities (paying a toll, being late for work, etc.). Thus, a toll may influence the complete daily schedule of a person, and not only the period when the toll is charged.

1

Our approach uses multi-agent simulations to model and simulate full daily plans. The agents are generated as a synthetic version of the real population living in the area of interest, matching as many demographic attributes as are available. Every agent then plans his/her day. These plans are simultaneously executed in the synthetic world, and the interaction between the agents (e.g. congestion) is computed. Agents are then given a chance to modify their plans, the plans are executed again, etc., that is, the agents adapt (co-evolution). Since the agents are virtual representations of real people, and since they move in a virtual representation of the real world, aspects of reality such as those discussed in the previous paragraph can be included in a conceptually straightforward way.

Each of our agents has a (possibly individual) utility function, based on the agent's performance during plan execution (i.e. including interaction). This utility function plays the same role as the fitness function in co-evolutionary genetic algorithms: Every agent attempts to improve his/her fitness, which, however, is based on the behaviour of the other agents. In consequence, changes in an agent's utility in reaction to a policy change can be directly interpreted as if and how much the agent benefits from the policy change, which allows to identify winners and losers of a policy measure. Since agents can adapt, their utility change reflects *indirect* utility: the utility change after agents have optimally adapted to the new circumstances. This adaptation takes place along those *choice dimensions* that are allowed by the simulation system, thus allowing fine-grained control over the adaptive reactions that are to be included into the scenario. Since utility is based on the performance in the synthetic world, it is conceptually straightforward to include constraints such as opening times: If an agent remains at a facility while the facility is closed, no utility will be accumulated during that time, and thus the agent will most probably search for a different time structure for his/her day.

As the change of the utility of an agent is the change of her/his personal performance, the (weighted) sum of all utility changes in reaction to a policy change is, in consequence, the change in the system's welfare. That is, the input of "user benefits" to a cost-benefit analysis can be taken directly from the simulation, and is therefore automatically consistent with the assumptions about the agents' optimization capabilities. This is in stark contrast to the traditional approach, where there is one model in which the agents react to a policy change, and a second model that assigns economic valuations to those changes. That approach completely separates the agents' optimization behaviour from the measurement of the economic benefit, thus making it difficult if not impossible to keep those consistent.

It might be worth emphasizing that the approach described in this paper, is meant to evaluate/appraise approaches that are designed by other means. There are other approaches, such as by Markose et al (2007) or Zhang and Levinson (2007), that concentrate on finding optimal designs, using agent-based approaches on the supply side. The emphasis of the present paper, in contrast, is on the travellers' behavioural realism, which we intend to improve much further in the near future. Ultimately, however, it is probable that the two approaches, supply-side agents on the one hand, and improved behavioural realism on the demand side, will merge.

This paper describes the agent-based approach in more detail. Section 2 describes the overall approach, concentrating on conceptual aspects, the co-evolutionary adaptation, and the scoring. Section 3 then describes a specific scenario, related to an illustrative study using data from the Zurich metropolitan area. The scenario consists of the geographic and socio-demographic input data, the computational modules, the toll scheme, and the specific simulation runs that were undertaken. In principle, the computational details should be rather unimportant, since agents are assumed to adapt optimally in the choice dimensions that are included in the simulation. In practice, however, at this point little is known about the robustness of these results, and for that reason more details are given. Section 4 describes the results of the toll simulation, in particular the reaction patterns that emerge. An important result is that agents react also outside tolled times – as they should. Section 5, finally, comes back to the economic analysis that was shortly sketched in the previous paragraphs.

2. SIMULATION STRUCTURE

2.1   Overview

Our simulation is constructed around 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 receives a 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 plan generation and traffic flow simulation. The system remembers several plans per agent, and utilities the performance of each plan. Agents normally chose the plan with the highest score, sometimes re-evaluate plans with bad utilities to test if it now performs better, and sometimes obtain new plans. For the purposes of this paper, a "score" is the same as a "utility". Further details will be given below.

A plan contains the itinerary of activities the agent wants to perform during the day, plus the intervening trips the agent must travel between activities (see Figure 1). 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 trip. This paper concentrates on "home" and "work" as the only activities, and "car" as the only mode. Our implementation supports additional activity types (see, e.g., Meister et al 2006) and additional modes of transport, but more time is needed to validate results given 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 incremental process, the input to a given module is a (possibly incomplete) plan, and the output is a plan with one or more decisions added. This paper will make use of two modules only: "activity time generator" and "router". Other modules are under investigation, and will be the topic of future publications. Once all agents' plans have been constructed, they can be fed into the traffic flow simulation. The traffic flow simulation 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.

```
<plan>
   <act type="home" link="123" end_time="06:30:00" />
   <leg mode="car">
      <route> 3 6 19 </route>
   </leg>
   <act type="work" link="456" duration="08:00:00" />
   …
</plan>
```

Figure 1: A "plan" in XML format. For the purposes of the present paper, a "leg" is the same as a "trip".

The outcome of the traffic flow simulation (e.g. congestion) depends on the planning decisions made by the agents and their 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 of the results 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 iteration cycles coupled with the agent database enable the agents to learn how to improve their plans over many iterations.

2.2   Co-evolutionary adaptation

The iterative approach outlined above describes co-evolutionary adaptation: Every agent attempts to improve, while everybody around him/her is doing the same. Several conceptual issues are related to this concept:

- If the adaptation ever came to an end (to a fixed point), and if the whole system were deterministic, then the system would be in a Nash equilibrium: No agent could improve by unilaterally doing something else. This notion is important, since it relates our work to standard game theory. There is, however, no guarantee that the Nash equilibrium is unique; i.e., depending on the initial conditions and on the exact nature of the adaptation process, different fixed points might be reached (Hofbauer and Sigmund 1998, Watling 1999). Clearly, the possibility of multiple outcomes is important in policy contexts.

- There is not even a need that such a system goes to a fixed point at all. Other possibilities are cycles or chaotic attractors (Hofbauer and Sigmund 1998, Watling 1999).

- Our system, however, is not deterministic. Although this is probably more realistic, and makes the system more robust, it makes the interpretation more complicated. One could argue that the system is ergodic (roughly: it eventually visits all possible system states) and should thus go to a steady state phase space density (Cantarella and Cascetta 1995). Yet, the phase space is huge and our typical iteration numbers comparatively small, and thus an instance of pseudo-stability (Watling 1996) or "broken" ergodicity (Palmer 1989) is much more probable. This means that the iterations might find a situation that is (apart from fluctuations) locally stable, but there is no guarantee that running the iterations longer might not lead to completely different solutions. In this paper, such a situation will be called "relaxed".

The uniqueness of the Nash equilibrium solution of traditional transport modeling was one of its most important features: Although static assignment does not necessarily have much to do with reality especially in heavily congested situations, at least it allows a conceptually sound comparison of different policy scenarios. Multi-agent simulations do not possess this uniqueness property, as can be shown already with simple but plausible examples (Daganzo 1998). It is as of yet unclear what the agent-based models will be able to put forward instead. One possibility is to not only consider the relaxed state, but also to pay closer attention to the realism of the initial condition and the adaptation behavior itself. Such models, then, would attempt to predict the full dynamic trajectory of the learning behavior of the simulated population when faced with a policy or infrastructure change.

2.3   Scores (= utilities) of plans

In order for adaptation to work in a meaningful way, it is necessary to be able to compare the performance of different plans. This is easiest achieved by assigning scores to plans. This is the same as the fitness function in genetic algorithms, or the objective function in optimization problems. Note once more that every agent has its own scoring function, and attempts to optimize for her-/himself.

In principle, arbitrary scoring schemes can be used (e.g. prospect theory, Avineri and Prashker 2003). Much of the approach, in fact, would still function even if agents were only able to rank plans at any given instance. In this work, a utility-based approach is used, and therefore the word "score" is replaced by "utility" in the following. 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 utility 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, as we consider the stay at home before the first departure; $U_{perf,i}$ 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 (and in the absence of constraints) all activities at their typical duration need to have the same marginal utility (since otherwise an agent could gain by expanding the activity with the largest marginal utility while shrinking the activity with the smallest marginal utility). β_{perf} is, in fact, the (marginal) opportunity cost of time, since it is the utility that is foregone when "nothing" is done (also see below the discussion re "coming early/waiting").

$t_{0,i}$ 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 $t_{0,i} = t_{*,i} \cdot \exp(-\zeta / t_{*,i})$ where ζ is a scaling constant set to 10 hours. With this specific form, $U_{perf,i}(t_{*,i}) = \beta_{perf} \cdot \zeta$ 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 €/h; usually negative) for being late, and $t_{late,i}$ is the duration of being late to activity j.

- 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 €/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 $-|\beta_{travel}| - \beta_{perf}$. In contrast, 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 values of the parameters of the Vickrey model (Vickrey 1969, Arnott et al 1993) if our simulation would just look for late arrival.

2.4 Discussion of the utility function

The above scoring function seems to be rather complicated, and seems to have many parameters. In the investigations presented in this paper, however, activities cannot be dropped, and therefore several aspects become unimportant. Take, for example, a person with a daily plan that consists of "home" and "work" only, with typical durations of $t_{*,home} = 16h$ and $t_{*,work} = 8h$. When home and work are at the

same locations, the person can be "at home" for 16 hours, and work for 8 hours, and that is the optimal solution, since the marginal utilities for both activities at these durations are equal. Note that the marginal utilities need to be decreasing (i.e. utilities need to be concave), because otherwise that optimum would not be unique and stable.

When home and work are not at the same location, the person needs to travel between the locations. The time for travel will eat with into the 16h at home and into the 8h at work. With small travel times, the reduction of these times will be approximately the same, since the marginal utilities are the same. With larger travel times, however, the curvature of the utility function will play a role, and the utility function with the stronger curvature will give up less time. A work duration that is strongly anchored near the 8h would, therefore, be modelled with a $t_{0,work}$ only slightly smaller than $t_{*,work}$, while a home duration this is *not* strongly anchored near the 16h would be modelled with a $t_{0,work}$ considerably smaller than $t_{*,work}$. This would lead to a result where the travel time would mostly be taken away from the time at home.

Overall, this implies the following for the investigation described in this paper:

- It is plausible to assume that utility functions are increasing and concave in the neighbourhood of the "typical durations" of the activities. Beyond these restrictions, the specific form of the function does *not* matter.

- As long as external constraints (e.g. work hour regulations, shop opening times) do not play an overarching role, marginal utilities at the typical durations should be the same at all activities. They might, however, vary from person to person, corresponding to different values of time.

- The curvature of the utility function at the typical duration is needed separately for each activity type.

This implies that any function that allows setting the slope and the curvature at a given "typical duration" should lead to similar results – typical candidate functions for this are log, power-laws, or second-degree polynomials. The parameters to be estimated are the following: (i) typical durations for each activity type; (ii) *one* marginal utility parameter for all activities, which also relates the utility of time to the utility of money; (iii) curvatures of the utility functions at the typical durations for each activity type. With respect to typical durations (i) and the opportunity cost of time (ii), one can have some confidence that they can be estimated from time use surveys (e.g. Stinson, 1999): Considerable progress has been made by Jara-Díaz and Guerra (2003) and Jara-Díaz et al (2007) concerning the methodology of estimating utility functions for complete days, including the income-dependency of the parameters. (iii) is more problematic, and the sensitivity of the results to the choice of the parameters needs to be systematically tested. Yet, given the abundance of time-use data, one can have some confidence that useful estimations will eventually be possible.

In addition, disutilities for travelling and for being late need to be derived. There is considerable work on disutilities for travelling, usually framed as "value of time" investigations. A problem is that, as explained above, two disutilities need to be added, one that is simply the opportunity cost of time, and one that is the *additional* disutility of travelling. A second problem is that estimations from discrete choice theory, abundant in travel behaviour research, measure *in*direct utility (e.g. Bates 2006), while what is needed for our simulations is *direct* utility. Once more, given the results of Jara-Díaz et al, one can have some confidence that these issues will eventually be resolved. With respect to the disutility of being late, the value used in our current investigations is ad-hoc, simply mirroring the famous value of Vickrey. There is, however, agreement in the transport community that especially for home-to-work trips, some kind of penalty for being late makes sense. The parameters can, in principle, be estimated from surveys concerning the willingness-to-pay for a more reliable transport system (e.g. Chen et al 2002).

3. SCENARIO

In this section, a specific scenario will be described in some detail. The purpose of this section is to illustrate the more general concepts described above, but also to provide some intuition as to what is feasible when dealing with real world scenarios.

3.1   Geographical area; population; initial demand

The scenario 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 2 (a)). It has the fairly typical size of 10 564 nodes and 28 624 links.

The simulated demand consists only of commuters that travel by car in the aforementioned region, resulting in 260 275 agents. For the results presented here, all travellers will have a simplified activity pattern of home-work-home. This is done for illustration and for verification purposes; future studies will contain full activity chains. 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.

3.2   Computational modules

A specific scenario always comes together with a specific selection of aspects that are included in the modeling, together with specific computational modules that are used. In the following sections, we describe the modules in more detail. Further research is needed in order to estimate how important the details of these modules are. Regarding the traffic flow simulation, which represents the physical reality, it should be possible to eventually come up with a calibrated and validated model that describes that physical reality in sufficient detail; and in fact, our own experience (unpublished) indicates that the results do *not* hinge critically on the selection of the traffic flow simulation. Yet, also the selection of the decision modules should, as long as they are able to scan the search space in a meaningful way, not be absolutely critical, since it is the scoring function described above that ultimately determines the agents' behavior. There is indeed some indication that simple mental modules, together with many iterations (= very long computing times), can lead to plausible results (Nagel et al 2004, Balmer et al 2005). Yet, the issue of computing times remains important (Meister et al 2006, Charypar et al 2006).

3.2.1   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 random "mutations" 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 naive, it is sufficient in order to obtain useful results. As stated above, 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 "plug-in", this module can be replaced by a more sophisticated implementation if desired. A more sophisticated scheduling module is already available (Meister et al 2006). After appropriate testing, this will be used in future studies.

3.2.2   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 essentially technical details, the implementation of such an algorithm is straightforward (Jacob et al 1999; Lefebvre and Balmer, 2007). 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 of departure time.

3.2.3   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-flow 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, at 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, e.g., will be used by the router), trip travel time, trip length, percentage of congestion, and so on.

3.2.4   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 assign each plan a "utility", as explained above.

It is important to note that our framework always uses actual plan performance for the utility. 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. Waddell et al 2003, 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 experienced (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.

1. The agent database starts with one complete plan per agent, which is marked as "selected".

2. The simulation executes these selected plans simultaneously and outputs events.

3. Each agent uses the events to calculate the utility of its selected plan and decides which plan to select for execution during the next iteration (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 utilities of the plans. The probabilities are of the form $p = e^{\beta \cdot U_j} / \sum_i e^{\beta \cdot U_i}$, where U_j is the utility of plan j, and β is an empirical constant. This is the familiar logit model (e.g. Ben-Akiva and Lerman 1985).

4. Next, the simulation executes the newly selected plans, that is, it goes back to 2.

This cycle 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 the long run, measures to stop the iterations automatically and consistently need to be found.

3.3   Tolls

To test the approach, two different road pricing policies were simulated: a "time toll" where the toll amount to pay depends on the time spent driving, and a "city toll" which is a distance-based charge for a certain area of Zurich. These will be explained in more detail in the following.
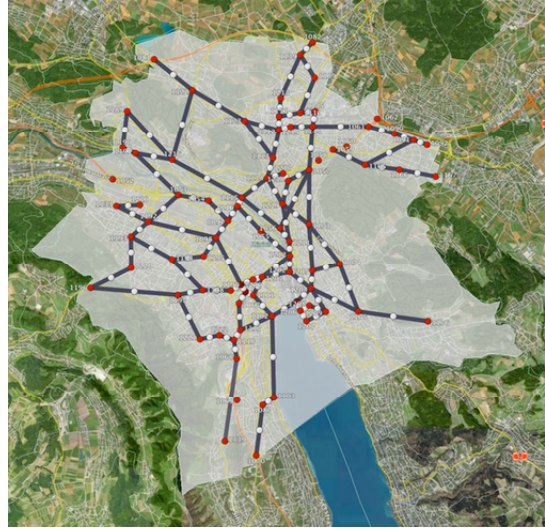
3.3.1   Time toll

The time-based policy is introduced for all travellers 24 hours a day, for the whole simulation area. For every hour spent on the road 12 € are charged; the computation of the toll is done in one-second increments. One could argue that this type of toll will never be implemented in the real world because it provides an incentive to speed. Yet, in-car devices for supervising speed limits are available (Garvill et al 2003) and could, in principle, be used together with such a toll. However, the main reason for using such a toll in a scientific study is that it approximates the time-dependent optimal toll in the bottleneck scenario analyzed by Arnott et al (1990): The time toll is more expensive for links where, and during time slots when, congestion is high. This becomes clear when one considers that the free speed travel time, and any toll levied on it, is foregone anyway, since in the set-up considered here the travellers can only adapt routes and times, i.e. they are forced by the simulation set-up to make the car trips under all circumstances. Therefore, only the time that comes on top of the free speed travel time is relevant for the decision-making of the agent – and that time is now made "more expensive" by the toll.

3.3.2   City toll

As a more realistic scenario, a hypothetical toll area was defined that covers the Zurich city area of administration, but not the motorways that lead into and partially around the city. The exclusion of the motorways is plausible because they are owned by the Swiss Confederation and not by the city of Zurich. This is a plausible scenario based on the status of the political discussion in Switzerland (Bundesrat, 2007). We are not aware of more specific politically discussed scenarios for Zurich; otherwise we would have used those. Figure 2 (b) shows the area and the tolled links. The diameter of the toll area is about 11 km. The toll is levied during morning rush hour, 7:15 am to 8:15 am, and is set to 1 €/km – in fact, toll levels of 0.5 €/km, 1 €/km, and 2 €/km were tested and out of those, the selected toll level yields the highest economic benefits (see below). The tolled 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 2: Scenario: Switzerland network with toll links for Zurich.

3.4   Simulation runs

The simulation is run for 100 iterations to retrieve a relaxed state in which the initial plans are adapted to the traffic conditions without any toll. Based on the hundredth iteration we simulate 200 iterations for each policy and for a base case where no policy is implemented. If not specified explicitly 10% of the agents adapt routes and 10% adapt activity times in each iteration. The following values for the parameters were used:

$\beta_{perf} = +6$ €/h, $\beta_{travel} = -6$ €/h, $\beta_{late} = -18$ €/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 $\beta_{perf} = +6$ €/h that it could accumulate instead. An agent that travels forgoes the same amount, plus a loss of 6 €/h for traveling. Finally, an agent that arrives late receives a penalty of 18 €/h.

The work activity's starting time is defined as 8:00 am for all agents. This is, once more, done to mirror the Vickrey scenario; more realistic values will be used in future studies.

The time toll is realized by increasing the penalty for traveling from -6 € per hour to -18 € per hour, i.e . $\beta_{travel} = -18$ €/h for the runs simulating the time toll.

4. RESULTS

In the following, we will present the results of three simulation runs comparing the base case with the results of the time- and distance toll. Thus, we will show which modifications in travel behavior are evoked when the policies are implemented in our scenario.

4.1   Effects of the time toll

The time toll reinforces the incentive of our agents to minimize travel time. On the other hand, there is still a strong incentive to start working at 08:00 am to avoid opportunity costs for waiting in case of early arrival or the penalty for arriving late. Making travel time more expensive means diminishing the influence of the schedule delay stimuli. Figure 3 shows the temporal distribution of the arrivals of agents at their activity locations for the base case and the toll scenarios. Comparing the base case curve and the time toll curve one can see that most of the agents still try to arrive at work at eight in

the morning. Some of them however are forced by the policy to arrive earlier than in the base case, i.e. more agents arrive between 06:00 am and approx. 7:30.
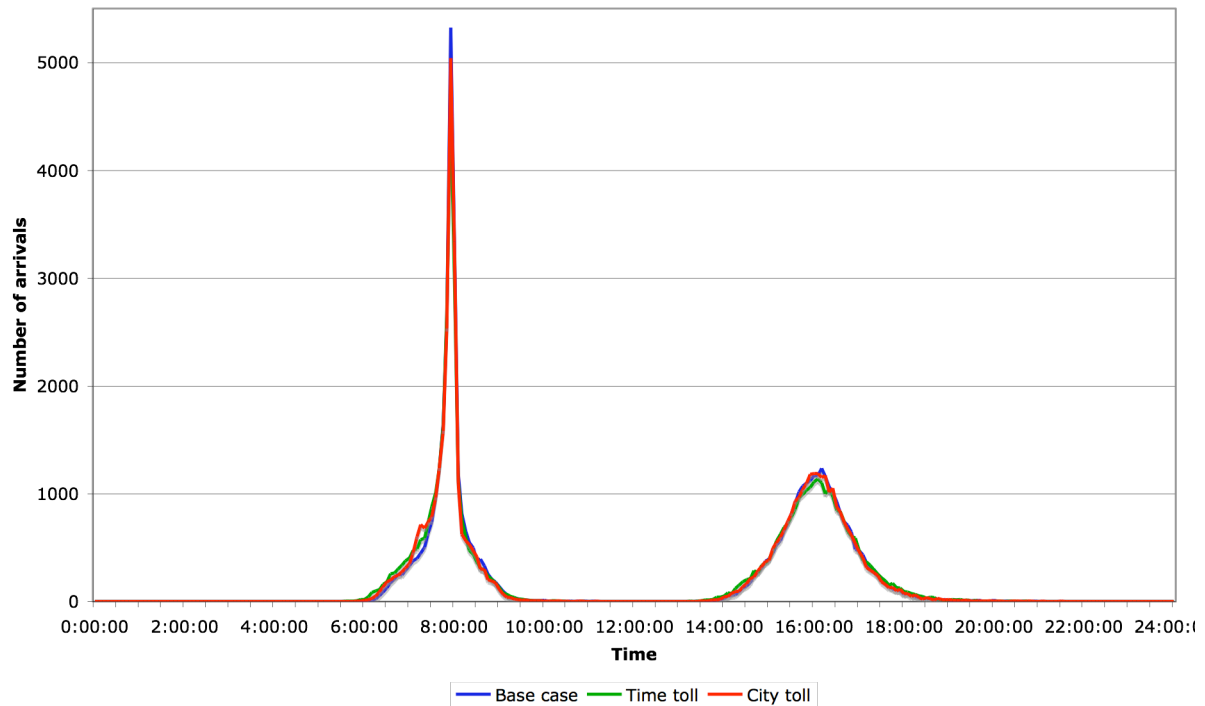

Figure 3: Number of arrivals of travelers over the time of day. The tolls force some agents to arrive earlier.

Figure 4 shows the number of travelers simultaneously on the road for the different scenarios. The area below the curves can be interpreted as the total time agents spend on the road. A smaller area means that people spend less time traveling. As some of the agents react to the toll by departing / arriving at different times or using other routes, the available infrastructure is used more efficiently, i.e. people have to spend less time traveling when the time toll is implemented.
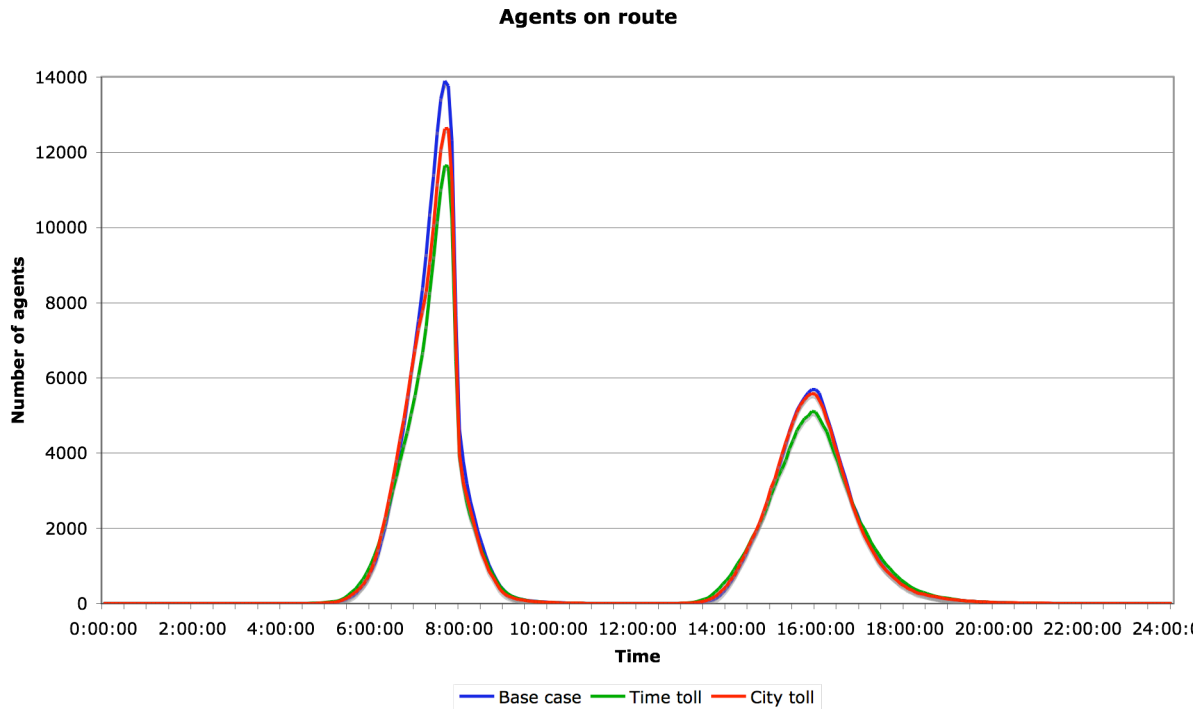
Agents on route

Figure 4: Number of travelers on the road over the time of day. With the tolls, the travel time is reduced.

Both results are as expected: The optimal toll in the Vickrey bottleneck scenario changes the departure pattern and reduces travel times, but it does not change the arrival pattern. In consequence, it makes sense that a toll that goes into the same direction has the same general effect. There are, however, also some differences: Figure 3 shows that the arrival pattern is not tri-angular, but is more strongly peaked at the desired arrival time at 8am. This is the result of the fact that departure and arrival locations are scattered throughout the region.

4.2   Effects of the city toll

As explained above, the city toll is a distance toll that is implemented only in the city area, and only between 7:15 and 8:15 am. Based on the arrival time distribution (Figure 3) and on the dynamics of the number of travellers on the road (Figure 4) one would say that its effect is in between the base case and the time toll. In addition, there is a slight toll time avoidance reaction in the morning. In contrast, there is virtually no effect in the evening. This makes sense since, as the morning arrival pattern is not much changed, there is little incentive to change anything in the evening when there is no city toll. This is in contrast to the time toll which is charged also in the evening.

4.3   Route adaptation only

The agent-based approach allows to selectively switch on and off certain choice dimensions of the travellers. In order to test this, additional runs were performed where the time adaptation of the agents in reaction to the tolls was switched off. This means that all departure times of the agents remain fixed also when the tolls are switched on – the agents can only react by moving to other routes. With the time toll, it is clear that there will be no reaction since the agents already use those routes that are fastest, given their departure time. But also with the city toll, Figure 5 implies that there is little temporal effect: The arrival time structure remains virtually unchanged, implying that the travel times remain unchanged and the toll thus does not relieve congestion. When interpreting these results, one should, however, keep in mind that in this set-up there is no elasticity of demand. This will be included into future studies and will certainly influence the results.
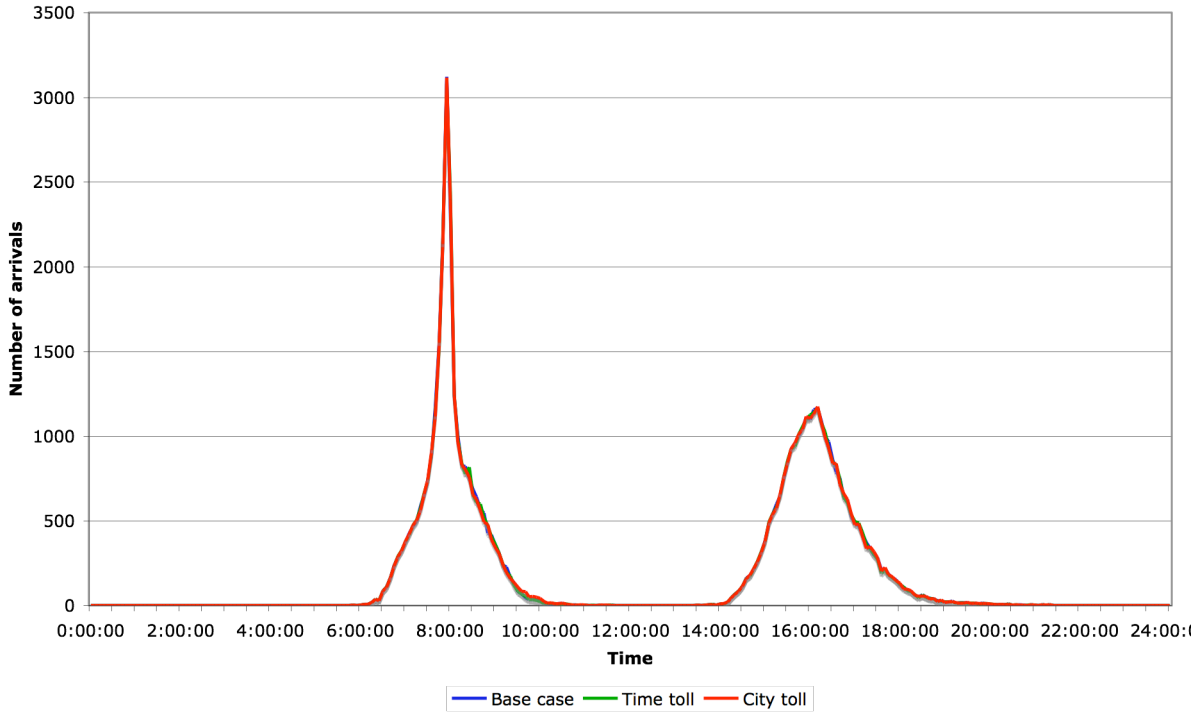
Figure 5: Number of arrivals of travelers over the time of day when no time adaptation is possible. The curves are nearly equal thus most reactions to the policies are time reactions.

4.4 Computational performance

The above runs need about 5 to 10 minutes of computing time per iteration; 100 iterations thus need about half a day. These numbers seems to be valid on any modern computer that we have tested, from laptops to workstations. Another element is memory. For the simulations for this paper, consisting of 10% of about 260,000 commuters, 2GB of memory are sufficient. A 100% simulation of the Zurich region, with full activity chains, takes about 20GB of memory.

5. ECONOMIC INTERPRETATION

5.1 Economic appraisal

Standard economic appraisal, as is for example used for cost-benefit-analysis (e.g. Pearce and Nash 1981), would now take the above traffic patterns as input, and attach economic valuations to them. One would, for example, for each link count the number of users and the average time they spent on the link. These numbers would then be compared between the base case and the scenario case. Typically, the travel time for a link would go down (say from t_1 to t_2) and the number of users would go up (say from n_1 to n_2). The economic gain consists of two contributions (e.g. Pearce and Nash 1981, Button 1993):

- Gains by existing users: $n_1 (t_1 – t_2)$, where the economic interpretation is that this was already the best option for those people before the modification of the system, and so the improvement of the system is fully counted for those people.

- Gains by new users: $(n_2 - n_1) (t_1 – t_2) / 2$, where the intuition is the following: New users one by one switch to the facility when the travel time is slowly decreased. A user that is switching is exactly neutral between two options. Any further improvement of the facility after the switch is then counted as benefit to that user. The first user that switches reaps nearly all the benefits from t_1 to t_2, while he last user that switches reaps nearly no benefits; this is overall approximated by $(t_1 – t_2)/2$ (known as the "rule of the half" in the literature).

13

The result of this procedure has a unit of time; it is then converted to monetary units by multiplying it with a value of time.

5.2   Identifying winners and losers in the multi-agent simulation

It should be clear, however, that this is too simplistic for the above scenario. In particular, the standard approach only counts travel time gains, but not schedule delay effects (caused by people choosing a non-preferred time of travel to avoid the toll). However, research indicates that schedule delay effects might contribute more than half of the economic effects of well-chosen time-dependent tolls (Arnott et al 1990). In addition, the standard approach is unable to look at equity aspects of the toll, since it is not able to differentiate between subgroups. Yet, equity effects, for example by age, income, gender, or race, are increasingly becoming important in any discussion of policy measures. Finally, the approach needs to assume a uniform value of time for all travellers, since the approach is not able to differentiate between trip purposes. Yet, it is well known that the value of time may differ by a substantial factor between leisure trips and business trips (for relevant Swiss results see Axhausen et al (forthcoming), which also differentiates by the length of the trip; see Jara-Diaz and Guevara (2003) for the problem of an equity-adjusted value of time).

Fortunately, with our multi-agent simulation, it is, in fact, not necessary to add the economic appraisal as it is conventionally done, since the agent utilities *already are the economic performance indicators of the system.* This is because the utility is the measure that every agent attempts to improve, and a higher utility directly measures the amount of improvement that an agent was able to reach. This automatically includes the schedule delay effects, since every agent will have optimally adjusted to any trade-off between time-dependent congestion, time-dependent toll, and schedule delay, including any personal restrictions that an agent may have, such as specific opening times. The multi-agent approach could also include different values of time, since they would be included as person-specific values of β_{perf}, β_{travel}, and β_{late}.

It is, thus, immediately possible to identify winners and losers of a policy. An example of such an analysis is Figure 6, which allocates the gains and losses after the introduction of the time toll to the residential locations of the agents. The analysis includes uniform redistribution of the toll revenues. In this case, there are both winners and losers close to Zurich, and mostly losers further away. The reason is that, far away from Zurich, only long distance travelers are included in the analysis, and these lose because the time toll punishes long distance travelers more, while the equal redistribution reimburses the same amount to everybody. Apart from this, there does not seem to be any structure, implying is that any structure that there may be is not geographically oriented.
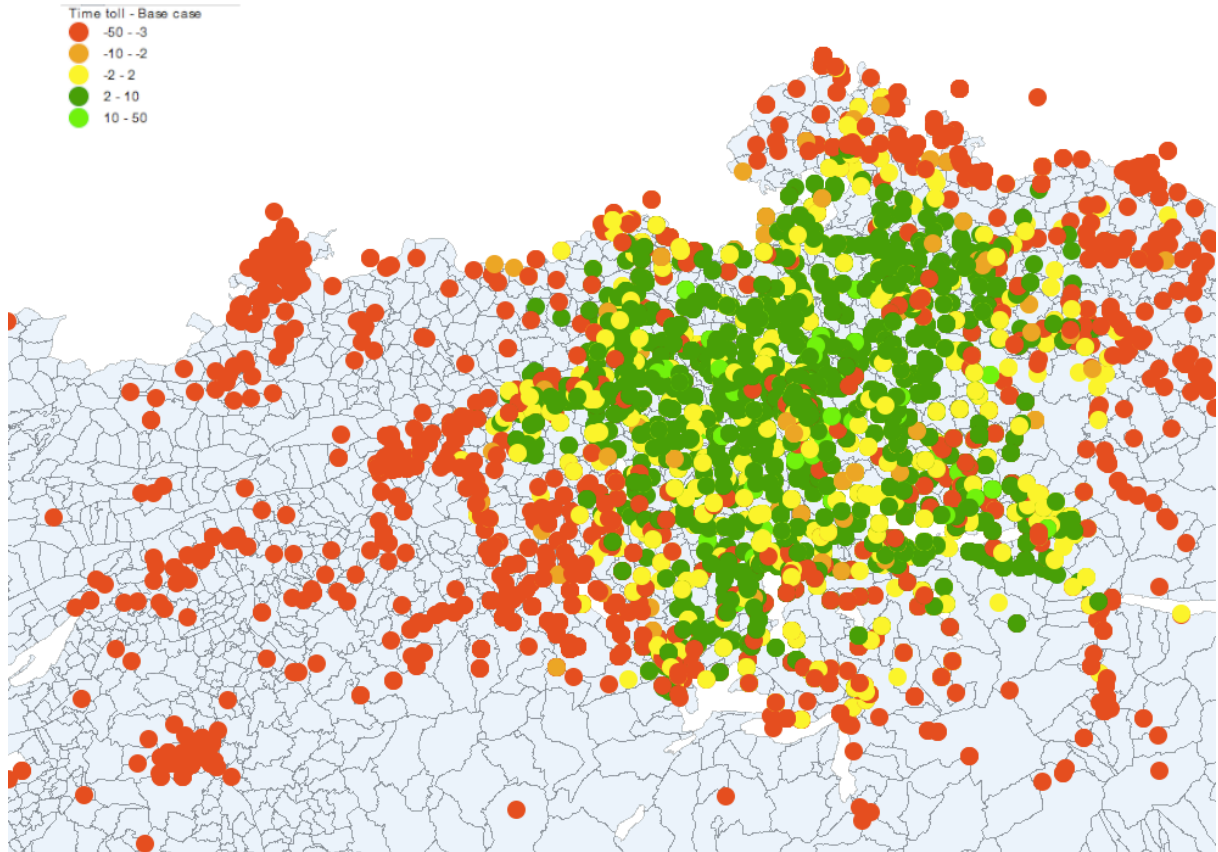
Figure 6: Spatial distribution of gains and losses of the time toll. Green: households with gains; red: households with losses. Households far away from Zurich lose because far away from Zurich only households with long distance trips are included, and these consistently lose with the time toll.

5.3 Aggregated economic benefits in the multi-agent simulation

For economic analysis, the gains and losses need to be aggregated. The simplest – utilitarian – way to do this is to sum up the utilities. The reasoning behind this would be that the monetarized utilities reflect willingness to pay of each individual for the change of the system by the policy (including the losers, which would need to receive money to accept), and thus the sum of these utilities reflects the aggregated willingness to pay. Table 1 shows one such analysis, for the scenario described above. On the left, one finds different entries for the base case. This is followed by two columns for the time toll, one depicting the new numbers after the policy introduction, the other the differences to the base case. This is followed by two columns of the same type for the city toll. For example, when switching on the time toll, the average travel time (per day per agent) decreases from 3832 to 3466 seconds. Average utility decreases by 10.18 €, including the utility loss from the toll payment. Since, however, average toll payments are 11.55 €, after redistribution of the toll revenues there is an average utility gain of 1.37 €. Multiplied by the number of agents this results in 357,476 € of utility gains.

The city toll reaches about ¾ of those gains (284,058 €), but with the advantages that already before redistribution there is a slight gain, and much less money needs to be moved around (less than 300,000 € instead of more than 3 million €) to obtain those gains.

When interpreting those results, one should keep in mind that, although our goal is to simulate realistic scenarios, the current example is still rather artificial: It is based on simplified home-work-home activities only, and there is no alternative to travelling by car (i.e. agents can neither decide to take an alternative mode nor to not travel at all). Future studies will remove these restrictions.

15

Table 1: Different contributions to the agents utilities for the different scenarios

	Base case		Time toll	Difference		City toll	Difference
Number of agents	260'890		260'890	0		260'890	0
Number of paying a-gents	0		260'890	260'890		57'990	57'990
Travel time (avg. per agent, sec.)	3'832		3'466	**-366**		3'682	**-150**
Utility (avg. per agent, €)	102.66		92.47	**-10.18**		102.69	**0.027**
Utilities (sum, €)	26'783'237		24'126'564	**-2'656'673**		26'790'317	**7'080**
Toll paid (avg. per pay-ing agent, €)	0.00		11.55	11.55		4.78	4.78
Toll paid (sum, €)	0		3'014'149	3'014'149		276'978	276'978
Utility after redistribu-tion of toll (avg. per agent, €)	102.66		104.03	**1.37**		103.75	**1.09**
Utility after toll redis-tribution (sum, €)	26'783'237		27'140'713	**357'476**		27'067'295	**284'058**

Table 2 shows what happens if the simulation allows only route adaptation in reaction to the toll. For the time toll, as expected the difference is essentially zero – the remaining difference can be attributed to the stochastic elements of the simulation. For the city toll, there are economic losses even after re-distribution of the toll (-136,476 € or -1.49 € per agent). This is due to the fact that the city toll will make people travel longer routes than without the toll. It is difficult to construct a toll for a city setting that generates economic benefits only from route adaptation. This also means, however, that one should be wary about toll analysis simulations where the only choice dimension is route adaptation – or in more general: As long as certain choice dimensions that are enacted in response to tolls are not included in the simulation models, the results may not be very useful.

Table 2: Different contributions to the utilities of the agents for the different scenarios if only route adaptation is permitted.

	Base case		Time toll	Difference		Distance toll	Difference
Number of agents	260'890		260'890	0		260'890	0
Number of paying agents	0		260'890	260'890		63'510	63'510
Travel time (avg. per agent, sec.)	3'638		3'630	**-8**		3'670	**32**
Utility (avg. per agent, €)	101.62		89.52	**-12.10**		100.13	**-1.49**
Utilities (sum, €)	26'511'958		23'354'332	**-3'157'626**		26'122'427	**-389'530**
Toll paid (avg. per agent, €)	0.00		12.10	12.10		3.98	3.98
Toll paid (sum, €)	0.00		3'156'769	3'156'769		253'054	253'054
Utility after redistribution of toll (avg. per agent, €)	101.62		101.62	**-0.003**		101.10	**-0.52**
Utility after toll redistribu-tion (sum, €)	26'511'958		26'511'101	**-857**		26'375'482	**-136'476**

An advantage of the agent-based approach is that the simulation and the appraisal are automatically consistent. Assume, for the purpose of illustration, that a toll is introduced for a fast facility, but that the facility can be circumvented by using a slower and untolled facility. Assume furthermore that the value of time is set to a small value in the simulation, but to a large one in the appraisal. In that situation, in the simulation much traffic would be diverted around the toll facility, since most travelers would rather spend more time than pay money. In the economic appraisal, however, that loss in time would be weighted very heavily. The result would be wrong, because if the travelers would use the larger value of time in the simulation, they would rather pay the toll and save time, thus making them better off. This example is arguably a bit trivial, but it is easy to design examples where such "incentive mismatches" can happen quite easily. In fact, the issue of the schedule delay, included in the simulation but excluded from current appraisal methods, is one example.

A similar consistency argument is made by de Jong et al (2005), where it is suggested to use the utility functions estimated for discrete choice models (or more precisely the expectation value, commonly called the "logsum term") directly for appraisal. This would have the same effect as our approach, i.e. that the model that is used for predicting the behavioral response is *the same* as the one that is used for evaluation/appraisal. This is undoubtedly a powerful approach, since it is by now established and common practice to estimate discrete choice models whereas realistic multi-agent simulations are still a challenge. The difference is that in our approach, the reactions of the travelers are directly and microscopically computed. The difference becomes clear when one, say, attempts to predict the effects of an area toll. In that situation, for the logsum approach it will be quite difficult to differentiate who will be affected by that toll and to what extent. One option would be to look at every person separately. Then, however, one ends up very close to a multi-agent simulation.

Finally, any aggregation by sub-groups is possible since the individual utilities are attached to every individual of the synthetic population, thus allowing filtering and aggregation by arbitrary criteria.

5.4   Discussion of the economic interpretation

Clearly, such an interpretation demands that the utility functions are not only correct for each individual agent, but that they can be compared between agents. The following arguments may be brought forward in this context:

- One could argue that it is unfair to just sum up the utilities of all agents, since this tends to, for example, weigh time gains/losses of the rich more heavily. As usual in this situation, it is always possible to weigh the utility gains of the agents before the summation according to whatever welfare function is to be optimized.

- One could argue that real people do not systematically optimize any scoring function at all, but rather do some kind of heuristic symbol processing to move through their days and lives (Simon 1997, Moss and Sent 1999). In that situation, one could replace the behavioral logic of our agents with an alternative logic. The simulation would still be able to predict future outcomes. The indicators, however, would now need to be based on other principles.

- The approach does not seem to include external effects. However, these could be included by computing the emissions caused by the traffic, and then compute the effects that these emissions have on the agents. Agents could even react to these effects, but that would be more relevant in the context of urban planning than in transport planning, e.g. because agents might move their residences to areas that are less impacted by emissions.

6.   CONCLUSION

It is shown that multi-agent simulations can be used to model people's reactions to policy measures in a conceptually straightforward and parsimonious way. In particular, when the simulated people (=agents) optimize individual utility functions, then changes in these utilities in reaction to a policy change directly measure individual gains or losses. This makes identifying winners or losers of a pol-

icy measure straightforward. Taking weighted sums over these individual utility changes then leads to aggregate welfare gains or losses. Since the aggregation is directly taken from the individuals, the weights can be based on individual demographic or geographic characteristics.

The approach is tested with a multi-agent simulation of traffic and travel behaviour of the Zurich metropolitan region in Switzerland. As policy, several tolling schemes are investigated. It is shown that the simulation can be used to model travelers' reactions to time-dependent tolls in a way most existing transportation planning tools are not able to do. As an example, it is demonstrated that route adjustment only, as is done by necessity in many traditional transport planning packages, results in no economic gains from the tolls – this is because the economic effects of tolls often stem from time reactions and from the elasticity of demand for car traffic. As time-dependent tolls are a much-debated subject in transportation politics, the ability to fully model such tolls and the reactions of travellers may help to find better toll schemes.

Overall, multi-agent simulations are able to approach a multitude of questions that current transportation tools are not able to answer. 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.

7. ACKNOWLEDGMENTS

8. BIBLIOGRAPHY

Avineri, E. and J. Prashker (2003) Sensitivity to uncertainty: Need for paradigm shift, *Transportation Research Record*, **1854**, 90-98.

Axhausen, K. (2008) Definition of movement and activity for transport modelling, in D. Hensher and K. Button (eds.), *Handbook of transport modelling*, 329–344, Elsevier, Oxford, 2nd edn.

Axhausen, K., S. Hess, A. König, G. Abay, J. Bates and M. Bierlaire (forthcoming) State of the art estimates of the Swiss value of travel time savings, *Transport Policy*.

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

Bates, J. (2006) Economic evaluation and transport modelling: Theory and practice, in K. Axhausen (ed.), *Moving through nets: The physical and social dimensions of travel*, chap. 10, 279–351, Elsevier.

Arnott, R., A. de Palma and R. Lindsey (1990) Economics of a bottleneck, *Journal of Urban Economics*, **27** (1) 111–130.

Ben-Akiva, M., M. Bierlaire, H. Koutsopoulos and R. Mishalani (1998) Dynamit: A simulation-based system for traffic prediction, DACCORS Short Term Forecasting Workshop, The Netherlands, February 1998. See also its.mit.edu.

Ben-Akiva, M. and S. R. Lerman (1985) *Discrete choice analysis*, The MIT Press, Cambridge, MA.

Bliemer, M. (2003) Analytical dynamic traffic assignment with interacting user-classes: Theoretical advances and applications using a variational inequality approach, Ph.D. Thesis, Technical University Delft, NL.

Bottom, J. (2000) Consistent anticipatory route guidance, Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Bundesrat (Government) of Switzerland (2007) Grünes Licht für Versuche mit Road Pricing, press conference, Dec/7 2007.

Button, K. (1993) *Transport economics*, 2nd edn., Edward Elgar Publishing Limited.

Cantarella, C. and E. Cascetta (1995) Dynamic process and equilibrium in transportation network: Towards a unifying theory, *Transportation Science A*, **25** (4) 305–329.

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

Chang, G., H. Mahmassani and R. Herman (1985) A macroparticle traffic simulation model to investigate peak-period commuter decision dynamics, *Transportation Research Record*, **1005**, 107–120.

Charypar, D., K. Axhausen and K. Nagel (2006) Implementing activity-based models: Accelerating the replanning process of agents using an evolution strategy, in *Proceedings of Swiss Transport Research Conference (STRC)*, Monte Verita, CH. See www.strc.ch.

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

Chen, A., Z. Ji and W. Recker (2002) Travel time reliability with risk-sensitive travelers, *Transportation Research Record*, **1783**, 27–33.

Daganzo, C. (1998) Queue spillovers in transportation networks with a route choice, *Transportation Science*, **32** (1) 3–11.

de Jong, G., A. Daly, M. Pieters and T. van der Hoorn (2005) The logsum as an evaluation measure: review of the literature and new results, in *presented at the 45th Congress of the European Regional Science Association*, Amsterdam.

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

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

Garvill, J., A. Marell and K. Westin (2003) Factors influencing drivers decision to install an electronic speed checker in the car, *Transportation Research Part F: Psychology and Behaviour*, **6** (1) 37–43.

Gawron, C. (1998) Simulation-based traffic assignment, Ph.D. Thesis, University of Cologne, Cologne, Germany, www.zaik.uni-koeln.de/AFS/publications/theses.html.

Hofbauer, J. and K. Sigmund (1998) *Evolutionary games and replicator dynamics*, Cambridge University Press.

Jacob, R. R., M. V. Marathe and K. Nagel (1999) A computational study of routing algorithms for realistic transportation networks, *ACM Journal of Experimental Algorithms*, **4** (1999es, Article No. 6).

Jara-Díaz, S. and C. Guevara (2003) Behind the subjective value of travel time savings, *Journal of Transport Economics and Policy*, **37** (1) 29–46.

Kaufman, D. E., K. E. Wunderlich and R. L. Smith (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 48109, May 1991.

Lefebvre, N. and M. Balmer (2007) Fast shortest path computation in time-dependent traffic networks, *Arbeitsbericht Verkehrs- und Raumplanung*, **439**, IVT, ETH Zürich, Zürich.

Markose, S., A. Alentorn, D. Koesrindartoto, P. Allen, P. Blythe and S. Grosso (2007) A smart market for passenger road transport (SMPRT) congestion: An application of computational mechanism design, *Journal of Economic Dynamics and Control*, **31** (6) 2001–2032.

MATSIM www page (accessed 2008) MultiAgent Transport SIMulation, www.matsim.org.

Meister, K., M. Balmer, K. Axhausen and K. Nagel (2006a) planomat: A comprehensive scheduler for a large-scale multi-agent transportation simulation, in *Proceedings of Swiss Transport Research Conference (STRC)*, Monte Verita, CH. See www.strc.ch.

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

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

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

Ortúzar, J. d. D. and L. Willumsen (1995) *Modelling transport*, Wiley, Chichester.

Palmer, R. (1989) Broken ergodicity, in D. L. Stein (ed.), *Lectures in the Sciences of Complexity*, vol. I of *Santa Fe Institute Studies in the Sciences of Complexity*, 275–300, Addison-Wesley, Redwood City.

Pearce, D. and C. Nash (1981) *Social appraisal of projects: A text in cost-benefit analysis*, Wiley & Sons, London.

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

Raney, B. and K. Nagel (2006) An improved framework for large-scale multi-agent simulations of travel behavior, in P. Rietveld, B. Jourquin and K. Westin (eds.), *Towards better performing European Transportation Systems*, 305–347, Routledge, London.

Simon, H. (1997) *An empirically based microeconomics*, Cambridge University Press.

Stinson, L. (1999) Measuring how people spend their time: a time-use survey design, *Monthly Labor Review*, 12–19.

Vickrey, W. S. (1969) Congestion theory and transport investment, *The American Economic Review*, **59** (2) 251–260.

Waddell, P., A. Borning, M. Noth, N. Freier, M. Becke and G. Ulfarsson (2003) Microsimulation of urban development and location choices: Design and implementation of UrbanSim, *Networks and Spatial Economics*, **3** (1) 43–67.

Watling, D. (1996) Asymmetric problems and stochastic process models of traffic assignment, *Transportation Research B*, **30** (5) 339–357.

Watling, D. (1999) Stability of the stochastic equilibrium assignment problem: a dynamical systems approach, *Transportation Research B*, **33**, 281–312.

Zhang, L. and D. Levinson (2007) Agent-based model of price competition, capacity choice, and product differentiation on congested networks, in *World Conference on Transport Research*, Berkeley, CA.