

Agent-based model for continuous activity planning with an open planning horizon

Journal Article**Author(s):**

Märki, Fabian; Charypar, David; Axhausen, Kay W. 

Publication date:

2014-07

Permanent link:

<https://doi.org/10.3929/ethz-b-000080319>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

Transportation 41(4), <https://doi.org/10.1007/s11116-014-9512-y>

Agent-based model for continuous activity planning with an open planning horizon

Fabian Märki · David Charypar · Kay W. Axhausen

Published online: 15 February 2014
© Springer Science+Business Media New York 2014

Abstract The paper proposes the microscopic travel demand model continuous target-based activity planning (C-TAP) that generates multi-week schedules by means of a continuous planning approach with an open planning horizon. C-TAP introduces behavioral targets to describe people's motivation to perform activities, and it uses a planning heuristic to make on-the-fly decisions about upcoming activities. The planning heuristic bases its decisions on three aspects: a discomfort index derived from deviations from agents' past performance with regard to their behavioral targets; the effectiveness of the immediate execution; and activity execution options available in the near future. The paper reports the results of a test scenario based on an existing 6-week continuous travel diary and validates C-TAP by comparing simulation results with observed behavioral patterns along several dimensions (weekday similarities, weekday execution probabilities of activities, transition probabilities between activities, duration distributions of activities, frequency distributions of activities, execution interval distributions of activities and weekly travel probability distributions). The results show that C-TAP has the capability to reproduce observed behavior and the flexibility to introduces new behavioral patterns.

Keywords Continuous activity planning · Multi-week schedules · Multi-week diary validation · Microscopic travel demand simulation · Continuous target-based activity planning (C-TAP)

F. Märki (✉) · D. Charypar · K. W. Axhausen
IVT, ETH Zurich, 8093 Zurich, Switzerland
e-mail: fabian.maerki@ivt.baug.ethz.ch

D. Charypar
e-mail: charypar@ivt.baug.ethz.ch

K. W. Axhausen
e-mail: axhausen@ivt.baug.ethz.ch

Introduction

Microscopic travel demand simulation software uses a direct representation of virtual people (usually referred to as agents) to generate demand in the form of activity plans (see e.g. Axhausen 1990b; Smith et al. 1995; Bowman 1998; Arentze et al. 2000; Miller and Roorda 2003; Bhat et al. 2004; Schnittger and Zumkeller 2004; Balmer 2007). For instance, Balmer's (2007) model uses agents who choose between different daily schedules. Activities from these schedules are executed, and the simulation results are returned to the planning process, allowing the agents to improve their schedules based on improved estimates of their generalized costs. This replanning step is repeated until the simulation reaches a stochastic user equilibrium with consistent travel demand and travel cost (Nagel and Flötteröd 2009). Simulating agents individually leads to high computational complexity, which often results in computational performance and memory issues. Microscopic models typically introduce restrictive constraints to counter such issues. For instance, Balmer limited the maximum simulation horizon of standard size scenarios to a single day, making it difficult to investigate effects occurring over a period of days or weeks. Another limitation of Balmer's model is that agents must commit to a specific day plan, making it challenging to simulate unexpected events realistically (Charypar et al. 2009; Dobler et al. 2012). Including such flexible behavior requires a different simulation approach that is capable of modeling demand continuously, i.e. agents should be able to make decisions about upcoming activities on-the-fly and with an open planning horizon (see also the empirical insights from the work of Doherty 2005).

We propose the microscopic travel demand simulation continuous target-based activity planning (C-TAP) that utilizes behavioral targets to guide agents through their decision space. These targets are closely related to observed behavior, e.g. the execution frequency or time spent on a given activity, and can take account of exogenous effects like social and cultural norms. Our agents continuously track their performance and compare it to their behavioral targets using observation windows of different durations. Deviations from the desired behavior cause *discomfort*, which is conveyed to a planning heuristic that makes decisions about which future activities the agents should execute. This enables agents to react spontaneously to unexpected events. At the same time, it also reduces memory consumption because agents do not need to keep track of complete schedules.

The final goal of this work is to provide a travel demand simulation featuring customizable components to model people's behavior and their environment. The resulting model should allow for simulation periods of several weeks and hundred thousands of agents within a reasonable computing time. Consequently, the challenge is to find a balance between behavioral realism and computational demand. In this paper we will focus on the introduction of the model and its configuration in order to reproduce observed behavior.

We will first discuss the model and its behavioral targets. We will then introduce the planning heuristic with its key features and describe how it utilizes the target-based model to determine which future activities to execute. This will be followed by a section on model calibration and validation. We will conclude the paper with an outlook on future research.

Related work

Arentze and Timmermans (2006) introduced need-based theory and its associated model for activity generation (Arentze and Timmermans 2009) that assumes that the utility of activities is a function of needs. Whereas Arentze and Timmermans postulated that needs motivate people to execute activities, we assume that people form a conception of their

desired performance as they reported in the data collected by the Swiss Federal Statistical Office (BFS) (2006) as well as in other travel diaries (e.g. Axhausen et al. 2002; Schönfelder 2006; Axhausen et al. 2007). Peoples' conception of performance is transformed into targets. Consequently, targets are directly observable in the available data, and it is therefore possible to overcome the *hidden source* problem of the need-based model. Nonetheless, we kept C-TAP generic in the sense that the satisfaction of needs may still be one possible target in C-TAP. We incorporate Winston's (1982) suggestion to use time-dependent utilities for activities (see also Axhausen (1990a, pp 34–38) for a summary or Gliebe and Kim (2010) for recent work in this tradition) and introduce time-dependent effectiveness functions to describe the effectiveness of activities and locations with respect to discomfort reduction. In Märki et al. (2011) we presented a need-based model that was designed for continuous simulation. C-TAP drops the need-based approach and introduces measurements (which we refer to as *targets*) that are directly observable in existing long-duration diaries. We expect this shift of focus to produce a coherent model and to simplify its utilization for modelers.

Behavioral model

Agents, representing virtual people, are the central component of C-TAP. We assume that each agent has a motivation to execute activities. C-TAP represents this motivation through behavioral targets, describing agents internal conditions based on personal desires, preferences, and goals. Deviations from behavioral targets result in discomfort, which induces agents to take action against the deviation; higher deviations result in greater discomfort, which in turn leads to a greater urge to take action. Agents can reduce their discomfort by executing activities at different locations, and we assume that agents prefer activity/location pairs that provide more discomfort reduction. This is similar to the work of Arentze and Timmermans (2009), who proposed that activity utility is a function of need reduction. Agents interact with their environment when they execute activities. Effectiveness functions describe external conditions agents face during their interaction and specify the efficacy of activity execution at a specific location.

Targets

The core assumption of this research is that people are motivated to execute activities and that they have a conception of their motivation in terms of the performance they want to achieve. People specify this performance through behavioral *targets* and try to comply with them across *observation windows* of different durations. For instance, a person might wish to exercise for roughly two hours twice a week. This person might be satisfied with one hour of exercise once a week and state that up to 2.5 h of exercise three times a week would still be compatible with his or her weekly schedule. The targeted behavior could be transformed into targets as follows:

- The *percentage of time* target defines the amount of the time a person would like to spend on an activity within an *observation window*. In order to simplify the modeler's task, it is possible to specify the total execution duration, which is then internally converted to the *percentage of time* target. For the above example, the modeler would specify a *target value* of 2 h of exercise, a *bandwidth* of $\frac{0.5}{-1}$ h (the upper and lower bounds of the target value) and an observation window of 1 week (see Fig. 1a).

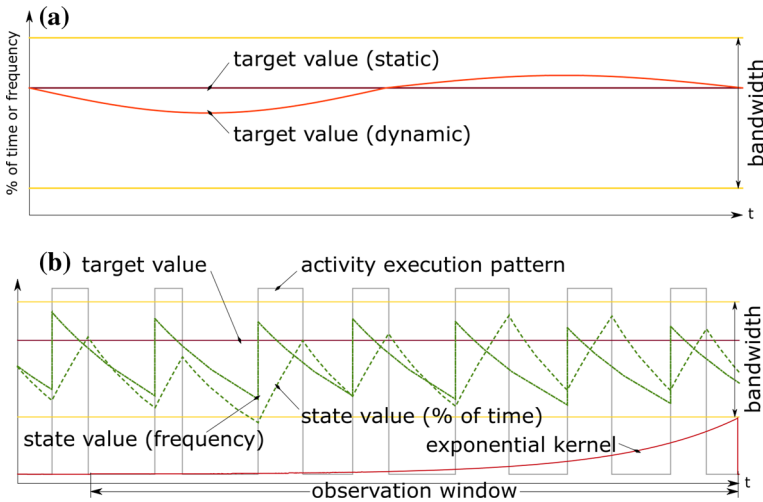


Fig. 1 Illustration of agent configuration and performance monitoring. **a** Schematic illustration of a target configuration that defines the average time a person would like to spend on executing an activity. *Target values* as well as upper and lower bounds (which define the behavioral *bandwidth*) can be static or dynamic. Static targets model constant behavioral patterns and dynamic targets model variable behavioral patterns (we refer to them as *behavioral rhythms*). Dynamic targets are modeled as functions in time. **b** Schematic illustration of performance monitoring. The *state value* (dashed green line) is calculated through a convolution of the *activity execution pattern* with an *exponential kernel*, resulting in an exponentially weighted moving average. The *observation window*, in which the person tries to comply with the target, defines the kernel length

- The *frequency* target defines the number of activity executions a person would like to accomplish within an observation window. For the above example, the modeler would specify a *target value* of two executions with a *bandwidth* of $\frac{+1}{-1}$ executions and an observation window of 1 week.

Playing tennis, swimming and hiking are all forms of exercise. At the same time, playing tennis is also a chance to socialize or even a means of maintaining business relationships (see e.g. Arentze and Timmermans 2009, for a discussion of multipurpose activities). C-TAP allows for activities to serve multiple targets, and it is possible to assign a target to several activities. This facilitates the configuration of interacting effects, as outlined above.

C-TAP supports static and dynamic targets. We use static targets to model constant behavioral patterns and dynamic targets to model variable behavioral patterns (we refer to them as *behavioral rhythms*). Behavioral rhythms are conditions in which activity executions vary across observation windows (e.g. leisure activities might be more common during weekends than weekdays).

Agents monitor their performance during the simulation and compare *state values* to their *target values* (see Fig. 1b). State values are exponentially discounted over the observation window of targets. This simulates a forgetting process in which agents give more weight to recent behavior and gradually discount their past performance.

Effectiveness functions

People seem to have time-dependent preferences for executing activities and/or visiting locations. Reasons for this behavior can be manifold, and they vary according to constraints (e.g. opening hours), norms (e.g. business hours) and contingencies (e.g. weather conditions) as well as combinations of such effects. Effectiveness functions are a simple but comprehensive way to describe such interdependencies and represent external conditions agents face when they interact with their environment. Effectiveness functions inform agents about the efficacy of activities and locations with respect to discomfort reduction and are expressed as the percentage of execution effectiveness. This is similar to the work of Winston (1982), who proposed time-dependent utilities for activities (see also Axhausen (1990a, pp 34–38) for a summary or Gliebe and Kim (2010) for recent work in this tradition). In C-TAP, effectiveness functions are a linear combination of different effects (constraints, norms etc.) and modelers can define different effectiveness functions for different locations (i.e. shops can have different opening hours). Effectiveness functions are agents' generic interface to the world and modelers can use them according to their modeling focus. Some effects that can be modeled by effectiveness functions include:

- *Shop opening hours for daily shopping activities* Agents can use this information either to determine whether they can execute a shopping activity without delay and for how long, or to estimate how long it will take until the next opportunity to engage in that activity will arise. Since effectiveness functions can be location dependent, it is also possible to model location dependent shop opening hours. Furthermore, the effectiveness function can also contain time dependent information about shop crowdedness. We assume that shopping in overcrowded shops is less efficient (smaller value) and therefore takes longer.
- *Daylight intensity for sleep activities* This function specifies the light intensity. Agents can, for instance, use this information as an indication of sleep effectiveness. We assume that people sleep at night and have already adapted to their current time zone.
- *Business hours for work activities* This function can be seen both as a cultural norm (different cultures may keep different business hours) and a social norm (social groups, e.g. professions, may have different business hours). Agents can use this information as an indication of work effectiveness. We assume that people depend on co-workers to be able to do their work (the degree of dependence can vary according to the profession).
- *Seasonal effects for leisure activities* This function is location dependent and combines different effects like the time of year and weather conditions. For example, a ski resort would be more effective during the winter months after a snowfall whereas a yacht club would be more effective during the summer months, when sunny weather and a good breeze are more likely. This function enables agents to follow seasonal rhythms by choosing to ski at a ski resort during the winter and to sail at a yacht club during the summer (see Märki et al. 2012, for an application of a similar example).

The decision model

Other approaches to agent-based microsimulations have exhibited various disadvantages. Balmer (2007) re-planned the same day until he produced an optimal state. This procedure led to high computational costs. Kuhnimhof and Gringmuth (2009) struggled with inflexibilities when agents should have spontaneously reacted to unexpected events.

Charypar and Nagel (2006) formulated the planning procedure as a reinforcement learning problem and reported that this approach performed poorly for large scenarios. We have introduced a decision heuristic to overcome the limitations described above. The proposed heuristic uses a continuous decision procedure, enabling agents to spontaneously react to unexpected events. Since a heuristic is intended to approximate a good solution, it is also possible to use incomplete knowledge about the state of mind and plans of other agents. This is helpful since complete knowledge generally leads to high computational and memory costs, especially if simulations are executed on distributed computation environments. One could argue that people seek optimal day plans and that applying a heuristic makes this infeasible. However, other authors (e.g. Simon 1955; Schlich 2004) doubt that behavior can be explained as a function of utility maximization that aims at a global optimum. They argue that people seem to maximize their utility but with a limited view similar to a search for a local optimum. One goal of this work is to demonstrate how far a decision procedure that approximates a good solution with limited information can reproduce real world behavior.

In the next section we will introduce the individual factors of the decision heuristic, which will be followed by a section outlining the decision procedure.

Individual factors of the decision heuristic

Discomfort

Discomfort builds on targets and is a function of the difference between targeted and actual behavior (i.e. the difference between *target values* and *state values*). Discomfort levels identify the urgency an agent experiences to take action against the origin of the discomfort. The discomfort an agent feels from an activity at time t is defined as

$$D(t) = \sum_{k=1}^n \left(f_{target}^k(t) - f_{state}^k(t) \right)^2 \cdot \begin{cases} w_1^k & \text{if } f_{state}^k(t) \leq f_{target}^k(t) \\ w_2^k & \text{otherwise} \end{cases} \tag{1}$$

$$w_1^k = \frac{1}{\left(f_{target}^k(t) - f_{lower-bound}^k(t) \right)^2} \tag{2}$$

$$w_2^k = \frac{1}{\left(f_{target}^k(t) - f_{upper-bound}^k(t) \right)^2} \tag{3}$$

the sum of the squared difference of the target value $f_{target}^k(t)$ and the state value $f_{state}^k(t)$ of all targets k (e.g. *frequency* and *percentage of time*) of the activity, normalized by the squared difference of the state value and the lower bound w_1^k if $f_{state}^k(t) \leq f_{target}^k(t)$, and otherwise by the squared difference of the state value and the upper bound w_2^k . The weight factors w_1^k and w_2^k normalize $D(t)$ and ensure that it equals to 1 at the lower and the upper bound.

Discomfort reduction

The reduction in discomfort that an agent experiences as a consequence of executing an activity at a specific location is defined as

$$DR(t_{es}, t_{ee}) = D(t_{es}) - D(t_{ee}) \tag{4}$$

which is the difference between the discomfort $D(t_{es})$ at the start of execution t_{es} and the expected discomfort $D(t_{ee})$ at the end of execution t_{ee} . The purpose of this difference is to introduce a preference for activity/location pairs that yield higher levels of discomfort reduction.

Execution effectiveness

The execution effectiveness an agent experiences for executing an activity at a given location is defined as

$$EE(t_{es}, t_{ee}) = \frac{\int_{t_{es}}^{t_{ee}} f_{effect}(t) dt}{t_{ee} - t_{es}} \tag{5}$$

whereby the integral of an effectiveness function $f_{effect}(t)$ (see “**Effectiveness functions**” section) between the start of execution t_{es} and the end of execution t_{ee} is normalized by the activity duration $t_{ee} - t_{es}$. The purpose of this index is to introduce a preference for executing activity/location pairs during efficient time windows (whereby efficiency is defined by whatever the effectiveness function represents, e.g. a combination of several effects like social and cultural norms), and it helps to prevent agents from executing activities during time windows when the agent cannot or can only partially execute the activity (e.g. because the shop closes).

Look-ahead index

Atkinson (1994) and Ioannou et al. (2001) highlighted the importance of information about future execution options for scheduling problems with time-window constraints. Effectiveness functions provide information about future execution options. For instance, shop opening hours inform agents about whether they can execute a shopping activity at a given time and for how long, or how long they must wait until the next opportunity to shop will arise. Agents can use such information to plan ahead, for instance by postponing activities with more execution options/higher prospective effectiveness in favor of more pressing activities for immediate execution (e.g. do *daily shopping* duties because shops are going to close soon, therefore postpone some other activities).

The aim of the look-ahead index is to provide agents with an awareness of potentially decreasing activity execution options. A simple approach for the example of the shop opening hours would be to measure the percentage of total time available for a potential execution of shopping in the near future (e.g. within the next 4 days). This approach has the disadvantage that it does not distinguish between execution options that will open soon and those that will open later. Applying higher weights to execution options that open sooner alleviates this problem. We do this through a convolution of the effectiveness function with an exponential kernel (see Fig. 2), similar to the convolution of state values (see Fig. 1b). The look-ahead value an agent receives for executing an activity at time t is defined as

$$LA(t) = \begin{cases} 1 + w_1 \cdot \left(1 - \int_0^h f_{effect}(t+x) \cdot kernel(x) dx \right) & \text{if } f_{effect}(t) > 0 \\ 1 & \text{otherwise} \end{cases} \tag{6}$$

where w_1 is a weight factor, h is the look-ahead horizon of the kernel, and $f_{effect}(t)$ the effectiveness function. Intuitively, one could understand $LA(t)$ as an index describing the urgency to execute an activity when prospective execution options decrease. Therefore, we




Fig. 2 Illustration of the look-ahead index for shop opening hours on short and long weekends, comparing different look-ahead windows (2 and 7 days). The higher the value the closer the end of the given shop opening hours window. The value is higher before weekends, indicating fewer shopping options in the near future. The look-ahead index with a kernel of 7 days can differentiate between short and long weekends (value is higher before a long weekend). Choosing the right kernel length is therefore important, so we propose a duration of approximately two to three times the average interval between the execution of two activities of the same type (e.g. $3 \cdot 2 \text{ days} = 6 \text{ days}$ for daily shopping)

designed it in such a way that it yields a value in the range of $[1..w_1]$ (1 for low urgency and a value approaching w_1 for decreasing execution options).

Execution time quota

The execution time quota an agent receives for executing an activity at a location is defined as

$$ETQ(t_{ts}, t_{es}, t_{ee}) = \frac{t_{ee} - t_{es}}{t_{ee} - t_{ts}} \tag{7}$$

which is the ratio between the activity duration $t_{ee} - t_{es}$ and the duration between the start of travel t_{ts} (defined by the end of the last activity at the previous location) and the end of execution t_{ee} . Accordingly, a higher share spent for traveling results in a smaller factor. In the current model, travel times are randomly drawn from a dynamic distribution, yielding higher expected travel times for increasing demand and travel distance. This parameter introduces an aversion to traveling and a preference for activity execution. Consequently, it produces a preference for accessible locations (locations that can be reached fast) and fosters activity chaining (see “Simulation results” section and Fig. 7).

The decision procedure

An agent decides upon his or her next action by applying a three-step decision procedure to all his or her promising activity/location pairs. This decision procedure combines different planning aspects into a heuristic function. These aspects include a comparison of the past performance with the targeted performance, an evaluation of the immediate execution options, an estimate of the prospective activity execution opportunities, a consideration of the travel time, and a random term modeling individual perception. The heuristic function HF is defined as

$$HF(t_{ts}, t_{es}, t_{ee}) = \underbrace{DR(t_{es}, t_{ee}) \cdot EE(t_{es}, t_{ee}) \cdot LA(t_{ee})}_{\text{step 1: activity duration}} \cdot \underbrace{ETQ(t_{ts}, t_{es}, t_{ee})}_{\text{step 2: gain per time invested}} \cdot (1 + \epsilon) \tag{8}$$

step 3: random fluctuation [optional]

which is the multiplication of the discomfort reduction DR , the execution effectiveness EE ,

the look-ahead index LA , the execution time quota ETQ , and a scaled random term $(1 + \epsilon)$ (see “[Individual factors of the decision heuristic](#)” section for an explanation of the factors). In a first step, an agent decides upon the duration of the activity by searching for the time of the end of execution t_{ee} , which maximizes the product of $DR \cdot EE \cdot LA$. This results in a preference for activity/location pairs which provide discomfort reduction, which are commonly executed (e.g. based on cultural or social norms), and which have a certain inflexibility with regard to being postponed to a later point in time. In a second step the agent multiplies the maximized value by ETQ , introducing a preference for accessible locations. So far, this decision procedure is consistent for all agents in the sense that agents having the same state and facing the same situation will make the same decision. Clearly, this does not match the heterogeneity of observed behavior (see e.g. Schönfelder 2006; Chikaraishi et al. 2010). A typical workaround to improve heterogeneity is to introduce an unexplained preference, modeled as a random variable altering the perception of alternatives (see e.g. McFadden and Train 2000; Horni et al. 2012, for other studies modeling heterogeneity in a population by random variables). Agents can incorporate a scaled random term into the decision-making procedure as an optional third step. This step activates agents’ *individual perceptions* and the decision procedure becomes nondeterministic. Here we have kept the decision-making procedure fully deterministic by disabling the third step. In a final step, the agent decides to implement the activity/location pair yielding the highest value of the heuristic.

Model validation

We validated C-TAP against an existing six-week continuous travel diary (Schönfelder 2006; Axhausen et al. 2007) carried out in the Canton of Thurgau, Switzerland in 2003. The survey comprises an activity record of 230 people from different social backgrounds and professions (e.g. employees, self-employed, retirees, homemakers, and students) living in 99 households. The validation demonstrates that C-TAP can be configured to reproduce observed behavior. It could then serve as a starting point to make forecasts based on changes in agents’ environment (modeled by effectiveness functions) or to make forecasts based on changes in agents’ behavioral targets (see also “[Discussion and conclusions](#)” section).

Model configuration

We distinguished between 11 activity types (home, work, work-related, education, socialize, daily shopping, long-term shopping, leisure active, leisure excursion, private business, short vacation) and implemented a model using 22 dynamic targets per agent (11 *percentage of time* and 11 *frequency* targets, we refer to dynamic targets as *behavioral rhythms*; see Fig. 1). These rhythms were automatically extracted from the survey through a convolution of the activity patterns with an exponential kernel of one week, resulting in an exponentially weighted moving average of six weeks (see Fig. 3). Weekly rhythms were then compiled by extracting the combined mean values of all exponentially weighted moving averages for the *target value* and the combined minimal and maximal values for the *bandwidth* (the upper and lower bounds of the target value). One could interpret such a rhythm as a behavioral path that an agent could take during the simulated weeks. Agents tend to stay in the middle of the path, but have the freedom to move toward the edges (the upper and lower bounds) resulting in increasing discomfort.





Fig. 3 Illustration of the extraction procedure of the weekly rhythm for the *percentage of time target of home*. The rhythm for the *frequency target of home* was extracted in the same manner and repeated for all remaining activities

Fig. 4 Illustration of the validation procedure. A model is configured based on the reported schedules of the people, and it is then used to simulate six consecutive weeks. Statistics are extracted from both observed and simulated schedules and then compared to give an understanding of C-TAP’s capability of reproducing behavior



Simulation results

We generate ten duplicates of each agent and initialize agents at random. Using additional agents introduces variability and copying existing agents preserves the sociodemographic distribution in the simulation. We simulated six consecutive weeks and validated the simulation results by comparing different statistics extracted from the survey and the simulation (see Fig. 4). These comparisons illustrate the behavior reproduction capability of C-TAP.

Figure 5 illustrates the average similarity of weekdays (e.g. Mondays compared to Tuesdays for the same person/agent) in the survey (first table) and in the simulation (second table). Similarities were measured by Joh’s (2004) multidimensional similarity-measurement function. Joh defined similarity as the effort necessary to realign a day’s activity sequence and activity timing into the sequence and timing of another day. Days are more similar the less effort is needed for this realignment. The first table shows that the same workdays (e.g. Thursday/Thursday) had the highest similarity (darkest color). This is followed by the similarity between workdays (e.g. Friday/Tuesday) and weekend days (Sunday/Saturday). Workdays compared to weekend days (e.g. Sunday/Thursday) showed the least similarity (lightest color). The simulation reproduced this similarity pattern with




Fig. 5 Illustration of day similarities based on Joh’s (2004) multidimensional similarity-measurement function with respect to activity sequence alignment and activity timing. **a** Survey, **b** simulation




Fig. 6 Illustration of the probabilities of the execution of activities at weekdays. A darker color indicates a higher probability for an activity to be executed on a given day. **a** Survey, **b** simulation

the exception of Mondays. This deviation may be explained by the activity *short vacation*. In the survey, people usually returned from a *short vacation* on Monday morning, resulting in dissimilarity in comparison to ordinary Mondays (leading to a lighter color for the survey). Due to the shorter execution duration of *short vacation* in the simulation (see Fig. 8a), agents returned earlier and spent Sunday night at home. This resulted in a Monday which was similar to ordinary Mondays (leading to a darker color for the simulation).

Figure 6 compares the probabilities of the execution of activities at weekdays. A darker color indicates a higher probability. People and agents executed the activities *education*, *work*, and *work-related* on workdays and in some cases on Saturdays. People and agents stayed at *home* on every day of the week with a tendency to stay away from *home* on weekends. *Leisure active* was executed on every day of the week, but, like *leisure excursion* and *socialize*, tended to occur more often on weekends. People and agents

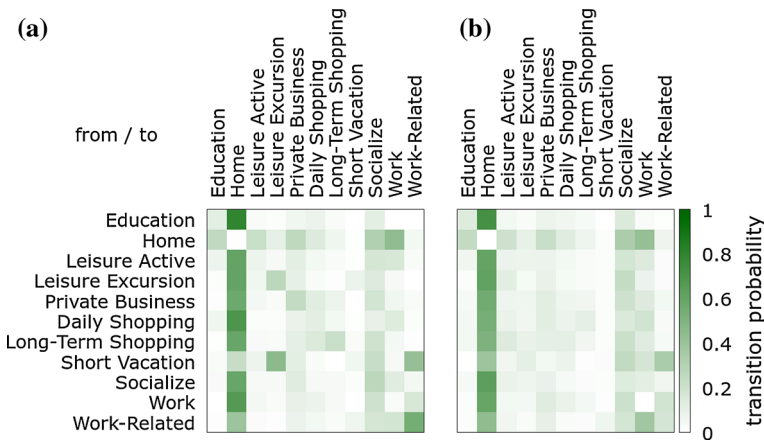


Fig. 7 Illustration of transition probabilities between activities. A darker color indicates a higher probability for a transition between given activities. **a** Survey, **b** simulation

performed *daily shopping* and *long-term shopping* on workdays and Saturdays. A generally slightly darker color of *leisure active* and *leisure excursion* in the simulation table indicates that these activities were less frequent in the survey (see also Fig. 8b). The simulation is apparently capable of reproducing the major execution patterns of the survey.

Figure 7 compares transition probabilities between activities. This defines the probability of an activity being executed right after another activity (e.g. home after work). A darker color indicates higher probability. People and agents were likely to go *home* after *work*, and there was a tendency to execute *work* after going *home*. Even though going *home* often succeeded other activities, it is possible to observe activity chaining in the survey and the simulation. *Socialize* neither had a clear predecessors in the survey nor the simulation. People re-executed activities (e.g. *leisure excursion*, *private business* and *work-related*). This also occurred in the simulation, but less distinctly. The simulation could develop new transitions (e.g. from *long-term shopping* to *leisure active*) and new tendencies could emerge (e.g. from *daily shopping* to *long-term shopping*). We see this as an indication that C-TAP has the flexibility to allow for new patterns. Based on these results we conclude that the simulation reproduced the major transition patterns of the survey.

Figure 8 compares the activity duration, frequency, and execution interval distributions. The figures group boxplots by activity type, forming pairs of boxplots to include survey (first) and simulation (second) results. The plots show that the target-based approach can represent different behavioral patterns, e.g. high frequency with long execution durations and short intervals (*home*), low frequency with short execution durations and long intervals (*long-term shopping*), or low frequency with long execution durations and long intervals (*short vacation*). Since behavioral targets build on weekly patterns (representing the average behavior in a 6-week survey), it is not surprising that C-TAP shows a tendency to produce less variability for durations and frequencies. The simulation reproduces most distributions with the exception of *short vacation* which deviates in mean value and observed variability. The survey does not provide many observations of this activity, and hence it is difficult to configure C-TAP automatically. Consequently, infrequent activities should be fine-tuned by hand using expert knowledge.

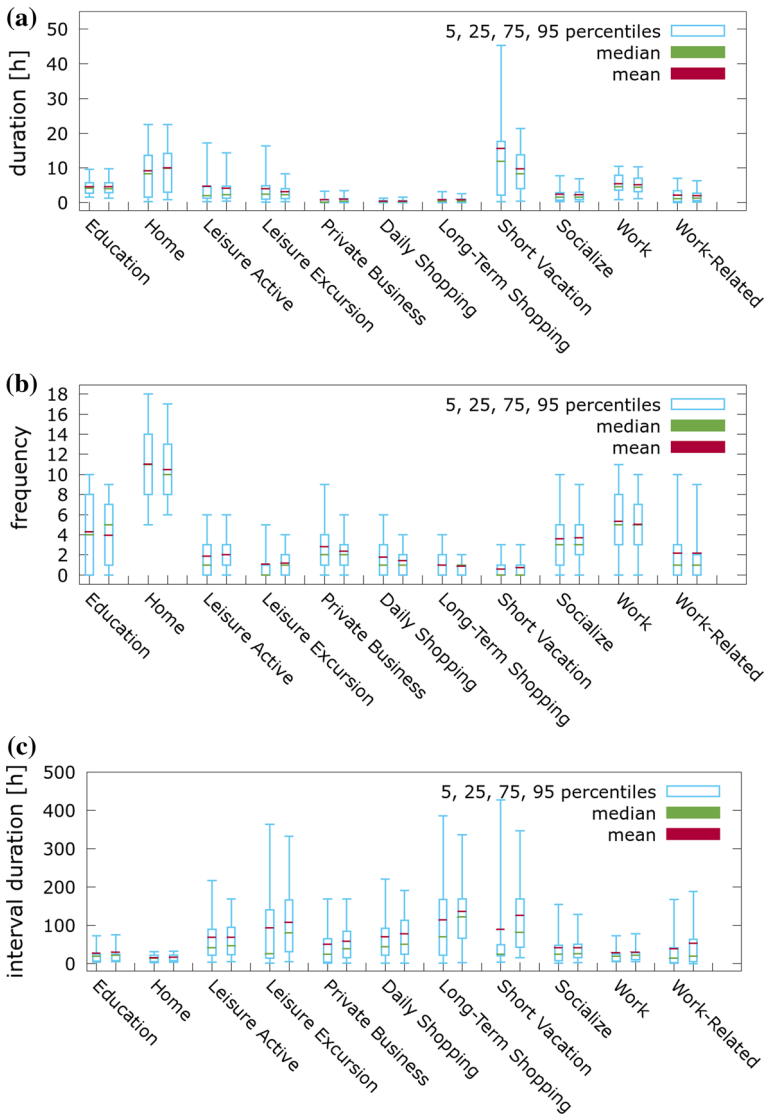


Fig. 8 Comparison of activity duration, frequency and execution interval distributions. Figures group boxplots by activity type, forming pairs of boxplots to include survey (first) and simulation (second) results. **a** Distributions of activity duration. **b** Distributions of activity frequency. **c** Distributions of activity execution interval

Figure 9 compares the travel behavior of people and agents. A higher value indicates a higher probability that traveling occurred at that time. Meeting these probabilities is important for a travel demand model and is an indication of its scheduling quality. The plot shows that the simulation reproduced major peaks, although some peaks are not as distinct (e.g. Wednesday morning), and the peak on Thursday morning occurred a bit later. In general, C-TAP shows a good capability for reproducing workday and weekend patterns.

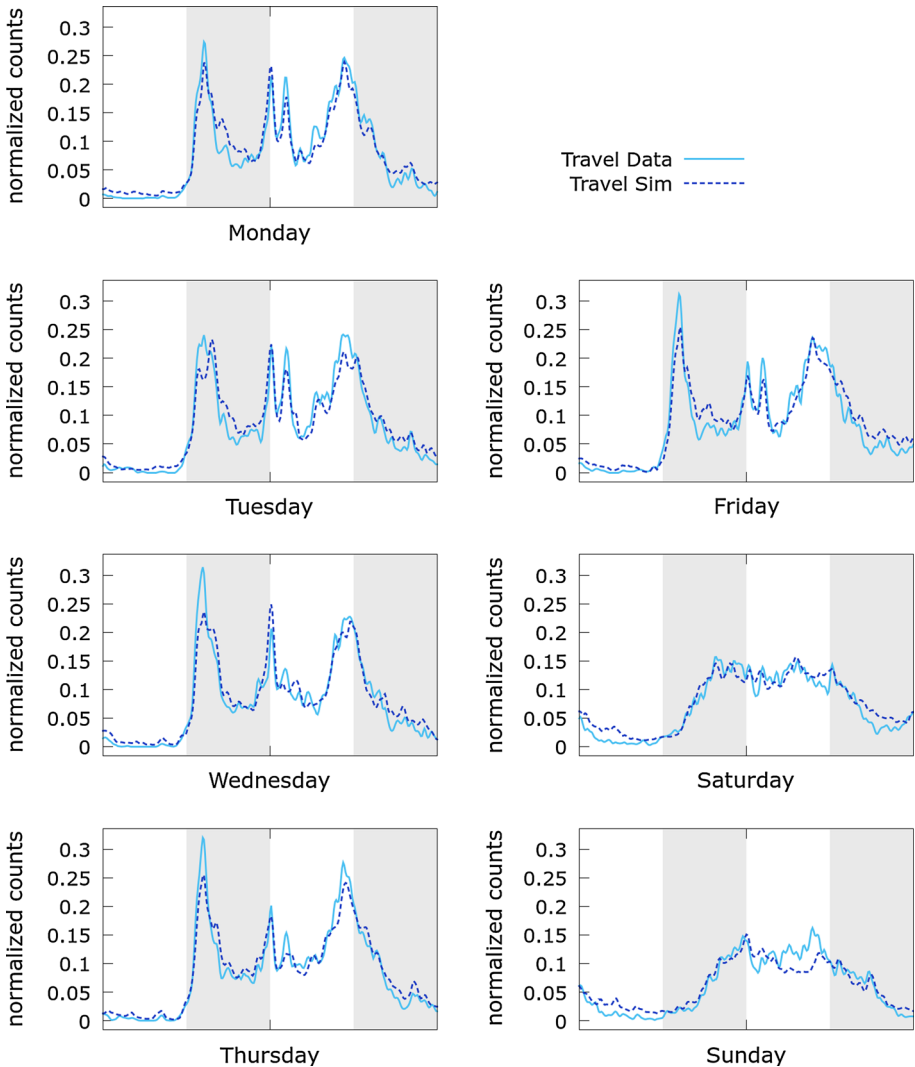


Fig. 9 Illustration of travel probability over the week. A higher value indicates a higher probability that travel occurs at that given time

Discussion and conclusions

Our first continuous model (Märki et al. 2011) used a need-based approach. It produced good results but made modelers' task difficult because the relationship between needs and activities was often unclear and necessitated modelers' interpretation. The target-based approach presented here introduces measurements that are directly observable in existing long-duration diaries. Consequently, these measurements can be extracted automatically (see “[Model configuration](#)” section), simplifying the modeling process. In case the available data does not provide enough samples (e.g. see *short vacation* in our validation), expert knowledge can still be used to improve simulation results.

Arentze and Timmermans (2007) pointed out that the dynamics of activity planning range from long-term dynamics related to life trajectories to short-term activity–travel rescheduling behavior, and several time horizons between these extremes. The decision procedure presented here plans activities on-the-fly considering different time horizons. On one hand, this procedure enables agents to react spontaneously to unexpected events since they can continuously adjust their decisions based on a changing environment. On the other hand, the kernel based approach of targets and look-ahead indices allows agents to consider different time horizons simultaneously. For instance, it would be possible to assign targets with kernel lengths of 1 week, 1 month, and 1 year to a work activity, ensuring that agents adhere to their weekly, monthly, and yearly work quota. In Märki et al. (2013), we showed how exogenous modules can induce agents to execute arbitrary activities. The same approach could be used to stimulate agents to execute activities like e.g., attending a funeral, even integrating activities nobody targets to execute into C-TAP.

The concept of time horizons limits agents' view into the past and future. Consequently, the decision procedure generates decision sequences based on limited information on agents' past performance and future opportunities, similar to a local optimization. Like other authors (e.g. Simon 1955; Schlich 2004), we believe this is behaviorally more authentic than the search for the globally best decision sequences (like e.g. Balmer 2007). Even though the decision procedure is still a simplification of reality it reproduced observed behavior convincingly and its computational demand enables simulation periods of several weeks and hundred thousands of agents within a reasonable time (the current algorithms allow a real time simulation of approximately one million agents using 16 threads on two Intel Xeon E7-8837 processors with eight-cores and 2.66 GHz each).

C-TAP derives activity durations from agents' state and external conditions. Agents monitor their performance during the simulation and compare it to their targeted behavior (see Fig. 1b). This allows for varying activity durations because they are based on the time needed to adjust *state values* to *target values* (i.e. reduce discomfort). Accordingly, an activity's execution duration varies with the time since its last execution. Furthermore, the execution duration also varies with its execution efficacy defined by the effectiveness function. E.g. if an effectiveness function takes shop crowdedness into account, shopping at an overcrowded shop takes longer than at a less crowded shop because the effectiveness influences the increase rate of the *state value*.

Many activities show a certain correlation between frequency and duration. C-TAP takes care of this by using *frequency* and *percentage of time* targets simultaneously. However, there are activities with a focus on frequency (e.g. escorting children to/from school) or on duration (e.g. reading a book). We model such activities by either dropping the minor target completely or by applying large bandwidths to the minor target, resulting in a reduced discomfort potential of the minor target in comparison to the major target.

C-TAP can account for heterogeneity in the population in several ways. First, it is possible to assign different effectiveness functions to the same activity/location pair of different agents. Accordingly, agents can perceive the same environment differently. Second, it is possible to assign different targets to agents. This is a way to introduce behavioral differences for individual agents or separate groups of agents. Third, it is also possible to introduce heterogeneity in the decision procedure by enabling *individual perception* (see “[The decision procedure](#)” section), modeled by the random term ϵ .

Modelers could use C-TAP by first generating a synthetic population, e.g. by using the algorithms proposed by Müller and Axhausen (2011, 2012). For each synthetic person they could then select the three to five most similar people from the six-week continuous travel diary used in this work (Schönfelder 2006; Axhausen et al. 2007). Similar to the presented

configuration procedure (see [Model configuration](#)“ section), modelers could then linearly combine the behavioral profiles of the selected people into weekly dynamic targets of an agent that represents the synthetic person. After the verification that the configured model reproduces observed behavior, modelers can start further investigations and make forecasts by manipulating targets and effectiveness functions to influence agents’ behavior. Manipulating targets is in accordance with changing agents’ behavior directly. C-TAP provides an interface that allows exogenous modules to manipulate agents behavior. Its application is manifold and ranges from modeling phase transitions in life (e.g. like the transition into marriage and parenthood) to modeling social interactions among groups of people (see Märki et al. 2013, for such an application). Manipulating effectiveness functions is in accordance with changing agents’ interaction environment. This could be used to test policy changes (e.g. to investigate possibilities to break morning peaks by starting classes later) or to model seasonal rhythms and weather conditions to see how agents react on seasonal differences (see Märki et al. 2012, for such an application).

In our research so far we have thoroughly and accurately tested C-TAP by visually inspecting several indices and distributions. In the future, it would be useful to have a quality metric that would allow an automatic inspection of individual behavioral aspects. Joh’s (2004) multidimensional similarity-measurement function gives an indication of the overall reproduction quality, but does not examine behavioral aspects separately. In addition to other distribution similarities (e.g. Bhattacharyya distance, Jeffrey divergence), we plan to use the *earth mover’s distance* proposed by Rubner et al. (2000) to automatically test model configurations for their reproduction capabilities of individual behavioral aspects.

Summary

This paper validates the microscopic travel demand simulation continuous target-based activity planning (C-TAP) that can continuously simulate agents’ behavior. The continuous nature of the simulation enables an investigation of traffic effects for durations of multiple weeks. Behavioral targets are central to C-TAP and represent agents’ internal state and describe their preferences and goals. These targets are closely related to statistical data provided by various sources, which simplifies the application of the demand simulation for modelers. Time-dependent effectiveness functions model various effects such as social and cultural norms and represent external conditions agents face when they interact with their environment. Agents keep track of their performance and compare it to behavioral targets. Deviations cause discomfort; this is conveyed to a planning heuristic that makes on-the-fly decisions about which upcoming activities agents should execute next. The required set of parameters was automatically extracted from an existing six-week continuous travel diary. The validation of the simulation by means of several groups of statistics and indices demonstrates that C-TAP reproduces various behavioral aspects observed in the data.

References

- Arentze, T.A., Hofman, F., Mourik, H., Timmermans, H.J.P.: Albatross: a multi-agent rule-based model of activity pattern decisions. *Transp. Res. Rec.* **1706**, 136–144 (2000)
- Arentze, T.A., Timmermans, H.J.P.: A new theory of dynamic activity generation. Paper presented at the 85th annual meeting of the Transportation Research Board, Washington, D.C., Jan 2006

- Arentze, T.A., Timmermans, H.J.P.: Modelling dynamics of activity–travel behaviour. Paper presented at the 12th international conference of Hong Kong Society for Transportation Studies, Hong Kong, Dec 2007
- Arentze, T.A., Timmermans, H.J.P.: A need-based model of multi-day, multi-person activity generation. *Transp. Res. Part B* **43**(2), 251–265 (2009)
- Atkinson, J.B.: A greedy look-ahead heuristic for combinatorial optimization: an application to vehicle scheduling with time windows. *J. Oper. Res. Soc.* **45**(6), 673–684 (1994)
- Axhausen, K.W.: Judging the day: a synthesis of the literature on measuring the utility of activity patterns. Working Paper, 561, Transport Studies Unit. University of Oxford, Oxford (1990a)
- Axhausen, K.W.: A simultaneous simulation of activity chains and traffic flow. In: Jones, P.M. (eds) *Developments in Dynamic and Activity-Based Approaches to Travel Analysis.*, pp. 206–225. Avebury, Aldershot (1990)
- Axhausen, K.W., Löchl, M., Schlich, R.: Fatigue in long duration surveys. *Transportation* **34**(2), 143–160 (2007)
- Axhausen, K.W., Zimmermann, A., Schönfelder, S., Rindsfuser, G., Haupt, T.: Observing the rhythms of daily life: a six-week travel diary. *Transportation* **29**(2), 95–124 (2002)
- Balmer, M.: Travel demand modeling for multi-agent traffic simulations: algorithms and systems. Ph.D. Thesis, ETH Zurich, Zurich, May 2007
- Bhat, C.R., Guo, J.Y., Srinivasan, S., Sivakumar, A.: A comprehensive econometric microsimulator for daily activity–travel patterns (CEMDAP). *Transp. Res. Rec.* **1894**, 57–66 (2004)
- Bowman, J.L.: The day activity schedule approach to travel demand analysis. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge (1998)
- Charypar, D., Horni, A., Axhausen, K.W.: Need-based activity planning in an agent-based environment. Paper presented at the 12th international conference on Travel Behaviour Research (IATBR), Jaipur, Dec 2009
- Charypar, D., Nagel, K.: Q-learning for flexible learning of daily activity plans. *Transp. Res. Rec.* **1935**, 163–169 (2006)
- Chikaraishi, M., Zhang, J., Fujiwara, A., Axhausen, K.W.: Exploring variation properties of time use behavior based on a multilevel multiple discrete-continuous extreme value model. *Transp. Res. Rec.* **2156**, 101–110 (2010)
- Dobler, C., Kowald, M., Schüssler, N., Axhausen, K.W.: Within-day replanning of exceptional events. Paper presented at the 91st annual meeting of the Transportation Research Board, Washington, D.C., January 2012
- Doherty, S.T.: How far in advance are activities planned?. Measurement challenges and analysis, *Transp. Res. Rec.* **1926**, 41–49 (2005)
- Gliebe, J.P., Kim, K.: Time-dependent utility in activity and travel choice behavior. *Transp. Res. Rec.* **2156**, 9–16 (2010)
- Horni, A., Nagel, K., Axhausen, K.W.: High-resolution destination choice in agent-based demand models. Paper presented at the 91st annual meeting of the Transportation Research Board, Washington, D.C., January 2012
- Ioannou, G., Kritikos, M., Prastacos, G.: A greedy look-ahead heuristic for the vehicle routing problem with time windows. *J. Oper. Res. Soc.* **52**(5), 523–537 (2001)
- Joh, C.-H.: Measuring and predicting adaptation in multidimensional activity–travel patterns. Ph.D. Thesis, Technical University Eindhoven, Eindhoven (2004)
- Kuhnimhof, T., Gringmuth, C.: Multiday multiagent model of travel behavior with activity scheduling. *Transp. Res. Rec.* **2134**, 178–185 (2009)
- Märki, F., Charypar, D., Axhausen, K.W.: Continuous activity planning for a continuous traffic simulation. *Transp. Res. Rec.* **2230**, 29–37 (2011)
- Märki, F., Charypar, D., Axhausen, K.W.: Location choice in a continuous model. Paper presented at the 13th international conference on Travel Behaviour Research (IATBR), Toronto, July 2012
- Märki, F., Charypar, D., Axhausen, K.W.: Integration of household interaction with a continuous simulation model. Paper presented at the 13th Swiss Transport Research Conference, Ascona, April 2013
- McFadden, D., Train, K.E.: Mixed MNL models for discrete response. *J. Appl. Econom.* **15**(5), 447–470 (2000)
- Miller, E.J., Roorda, M.J.: A prototype model of 24-h household activity scheduling for the Toronto area. *Transp. Res. Rec.* **1831**, 114–121 (2003)
- Müller, K., Axhausen, K.W.: Hierarchical IPF: Generating a synthetic population for Switzerland. Paper presented at the 51st congress of the European Regional Science Association, Barcelona, September 2011

- Müller, K., Axhausen, K.W.: Preparing the Swiss Public-Use Sample for generating a synthetic population of Switzerland. Paper presented at the 12th Swiss Transport Research Conference, Ascona, May 2012
- Nagel, K., Flötteröd, G.: Agent-based traffic assignment: Going from trips to behavioral travelers. Paper presented at the 12th international conference on Travel Behaviour Research (IATBR), Jaipur, December 2009
- Rubner, Y., Tomasi, C., Guibas, L.J.: The earth mover's distance as a metric for image retrieval. *Int. J. Comput. Vis.* **40**(2), 99–121 (2000)
- Schlich, R.: Verhaltenshomogene Gruppen in Längsschnitterhebungen. Ph.D. Thesis, ETH Zurich, Zurich, April 2004
- Schnittger, S., Zumkeller D.: Longitudinal microsimulation as a tool to merge transport planning and traffic engineering models: the MobiTopp model. Paper presented at the European Transport Conference, Strasbourg, Oct 2004
- Schönfelder, S.: Urban rhythms: modelling the rhythms of individual travel behaviour. Ph.D. Thesis, ETH Zurich, Zurich (2006)
- Simon, H.: A behavioral model of rational choice. *Quart. J. Econ.* **69**(1), 99–118 (1955)
- Smith, L., Beckman, R.J., Anson, D., Nagel, K., Williams, M.E.: TRANSIMS: transportation analysis and simulation system. Paper presented at the 5th TRB national transportation planning methods applications conference, Seattle, April 1995
- Swiss Federal Statistical Office (BFS): Ergebnisse des Mikrozensus 2005 zum Verkehrsverhalten, Swiss Federal Statistical Office (BFS), Neuchatel (2006)
- Winston, G.C.: *The Timing of Economic Activities*. Cambridge University Press, Cambridge (1982)

Fabian Märki is a Ph.D. candidate at the Institute for Transport Planning and Systems at ETH Zürich. His research interest is in agent-based algorithms for the efficient simulation of travel behavior in large scenarios. Fabian Märki holds a BS in computer science from the University of Applied Sciences Northwestern Switzerland and a MS with distinction in research from Stanford University. Prior to starting a Ph.D., Fabian Märki worked as a research fellow at the Institute of 4D Technologies of the University of Applied Sciences Northwestern Switzerland where he managed research projects and was involved in teaching computer science courses.

David Charypar is a research fellow at the Institute for Transport Planning and Systems at ETH Zürich. He studied computer science, and in his Ph.D. he focused on agent-based transport modeling, event-driven flow simulation, and parallelization. Thanks to a general attention to efficiency he contributed to making agent-based transport models applicable to large-scale scenarios. Currently, his research focuses on continuous simulation, parallelization, dynamic load balancing, and supercomputing.

K. W. Axhausen is Professor of Transport Planning at ETH Zürich. Before he worked at the Leopold-Franzens Universität Innsbruck, Imperial College London, the University of Oxford and the Universität Karlsruhe. He has been involved in the measurement and modelling of travel behaviour for the last 30 years contributing especially to the literature on stated preferences, micro-simulation of travel behaviour, valuation of travel time and its components, parking behaviour, activity scheduling and travel diary data collection. Current work focuses on the agent-based micro-simulation toolkit MATSim (see <http://www.matsim.org>) and on the land-use/transport interaction.