A Continuous Simulation Concept for Daily Travel

Fabian Märki

David Charypar

Kay W. Axhausen

STRC 2011 **May 2011**

STRC

11th Swiss Transport Research Conference
Monte Verità / Ascona, May 11 – 13, 2011

STRC 2011

A Continuous Simulation Concept for Daily Travel

Fabian Märki
IVT
ETH Zurich
CH-8093 Zurich
phone: +41-44-633 33 25
fax: +41-44-633 10 57
fabian.maerki@ivt.baug.ethz.ch

David Charypar
IVT
ETH Zurich
CH-8093 Zurich
phone: +41-44-633 35 62
fax: +41-44-633 10 57
charypar@ivt.baug.ethz.ch

Kay W. Axhausen
IVT
ETH Zurich
CH-8093 Zurich
phone: +41-44-633 39 43
fax: +41-44-633 10 57
axhausen@ivt.baug.ethz.ch

May 2011

Abstract

This paper proposes a microscopic traffic simulation that employs a continuous planning approach together with an open time horizon. It uses behavioral guidelines and the concept of *projects* to model people's motivation to execute activities. People's behavior originates from a planning heuristic making on the fly decisions about upcoming activities. The planning heuristic bases its decisions on the available activity options in the near planning future and on a discomfort measure which is derived from deviations between people's performance and their behavioral guidelines.

Keywords

behavioral guidelines, personal projects, continuous activity generation, open time horizon, microscopic traffic simulation

1 Introduction

Microscopic travel demand simulation softwares simulate each virtual person (referred to as agent) individually, often resulting in high computational complexity. Balmer (2007) uses agents which choose between different daily schedules. Activities of these schedules are executed and simulation results are handed back to the planning process, allowing 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)). This approach experiences computational performance issues, limiting its maximal simulation horizon (Charypar *et al.* (2009)). Another limitation is that agents must commit themselves to a specific daily schedule, making it difficult to simulate unexpected events realistically. As a consequence, a different simulation approach becomes necessary that is capable to model demand continuously with an open time horizon.

We propose a microscopic travel demand simulation that utilizes behavioral guidelines to represent agents' decision space. Guidelines can represent social and cultural norms and are closely related to observed behavior, simplifying model utilization for practitioners. We use the concept of projects to model non-recurring tasks. Projects influence behavioral guidelines during a specific time period. Agents continuously track their performance and compare it to their behavioral guidelines using observation windows of different temporal ranges. Deviations from the desired behavior cause discomfort which is conveyed to a planning heuristic that makes decisions about future activities 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 daily schedules.

The remainder of this paper is structured as follows: first, we discuss the model and its behavioral guidelines. We then introduce the concept of projects, describe how it influences behavioral guidelines and illustrate the concept with several examples. The next section describes the planning heuristic and its key features. We conclude the paper with a discussion of next steps towards implementation.

2 Other Work

Arentze and Timmermans (2006) introduced need-based theory and proposed a model for activity generation (Arentze and Timmermans (2009)) that assumes utilities of activities are a dynamic function of needs. Whereas Arentze and Timmermans used needs as people's motivation to execute activities, we assume that people have a direct perception of their motivation and describe their desired performance through behavioral guidelines. Axhausen (1998) and Schönfelder and Axhausen (2009) proposed projects as a coordinated set of activities, tied together by a common goal or outcome. Miller (2005) technically applied projects to organize complex human behavior. We see projects as a mechanism that temporally influences behavioral guidelines and use it to model non-recurring tasks. We pick up Gliebe and Kim (2010)'s suggestion to use time-dependent utilities and introduce time-dependent effectiveness guidelines.

3 Behavioral Model

Agents are the central component of our model and represent virtual people. Each agent has a motivation to do different things and specifies its desired performance through behavioral guidelines. Deviations to behavioral guidelines result in discomfort which induces agents to take action against the deviation; higher deviations result in higher discomfort which in turn leads to a higher urge to take action. Agents can reduce discomfort through the execution of activities at different locations. We assume that agents implement activity-location pairs that provide most discomfort reduction per time. This is similar to Arentze and Timmermans (2009), who propose activity utility as a function of need reduction.

The core assumption of this work is that people have a motivation to execute activities and that they have a perception of their motivation in form of a desired performance. People specify this performance through *behavioral guidelines* and try to comply with them across *observation windows* of different duration. For instance, a person would like to play 2^{+2}_{-1} hours of tennis about twice per week but not more than six times per month. In this example, the person specifies a *reference value* of 2 hours of tennis, a *bandwidth* of $^{+2}_{-1}$ hours, and two observation windows (per week and per month) in which the person tries to comply to the guideline.

3.1 Behavioral Guidelines

3.1.1 Observation Window Guidelines

Activity *execution frequency* and *cumulative execution duration* are two guidelines with observation windows. Both guidelines define a reference value (value a person tries to target) and a bandwidth (upper and lower bound of the reference value). For instance, a modeler could specify these guidelines for a sport activity as follows:

- **Frequency:** $2\times^{+2}_{-1}$ **per week.** The desired frequency is twice per week and the agent experiences a limited discomfort in the range of [1..4] times per week. The observation window is one week.

Figure 1: Illustration of discounted monitoring values

(a) Execution of an activity

(b) Discounted *percentage of time* spent for an activity and discounted *frequency* of execution

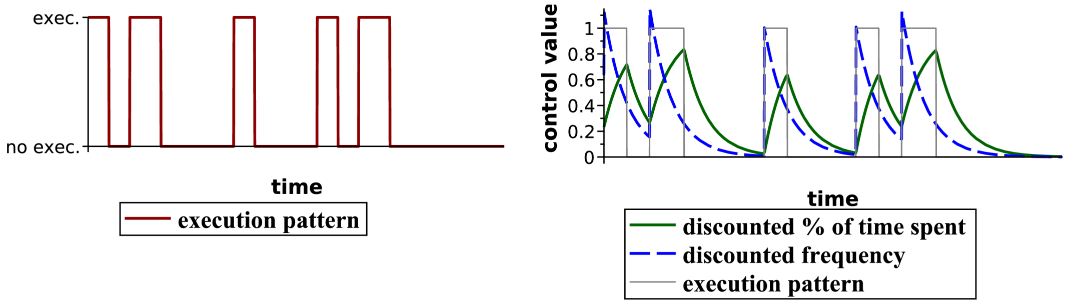

- **Cumulative duration:** $4h^{+2}_{-2}$ **per week.** The desired cumulative execution duration is $4h$ per week and the agent experiences a limited discomfort in the range of [2..6] hours per week. The observation window is one week.

Agents use *monitoring values* to record their performance and compare these values to guidelines. Monitoring values are exponentially discounted over their observation window. This is achieved by a convolution with an exponential kernel (see Fig. 1). Accordingly, agents give recent behavior more weight and *forget* their behavioral performance beyond the observation window. This is in contrast to a S-shaped functional form, as it is used by Arentze and Timmermans (2009), which does not facilitate forgetting of past performance. Internally, the simulation converts cumulative execution duration guidelines into *percentage of time* guidelines which define the percentage of total time an agent should spend for an activity. In our example, the reference value is $2 \cdot 4/(7 \cdot 24) = 4.76\%$ with an upper bound of $2 \cdot 6/(7 \cdot 24) = 7.14\%$ and a lower bound of $2 \cdot 2/(7 \cdot 24) = 2.38\%$.

3.1.2 Execution Duration Guideline

The execution duration guideline specifies how long an agent should spend for one activity execution with an upper and lower bound. For instance, a modeler could specify the satisfaction duration guideline for a sport activity as follows:

- **Duration:** $2h^{+2}_{-1}$ **per execution:** The reference duration is 2 hours per execution and the agent experiences a limited discomfort if it chooses a duration in the range of [1..4] hours.

3.1.3 Effectiveness Guideline

Similar to Gliebe and Kim (2010) who used time-dependent utilities, we introduce a time-dependent effectiveness guideline. This guideline informs agents about how effective it is to execute an activity at a specific time through a value in the range of $[0..1]$. The effectiveness guideline is a broad concept and can model different effects. Possible examples are:

- **Shop opening hours for a *daily shopping* activity.** This guideline takes the value of one when shops are open and zero whenever they are closed. Agents can use this information to either determine if they can shop and for how long or how long it takes until they can shop next time.
- **Daylight intensity for a *sleep* activity.** This guideline specifies the light intensity. Agents can use this information e.g. as an indication of sleep effectiveness. Hereby, we assume that people sleep at night and have already adapted to their current timezone.
- **Business hours for a *work* activity.** This guideline can be seen as a cultural norm (cultures may have different business hours) and a social norm (social groups, e.g. professions, may have different business hours). Agents can use this information e.g. as an indication of work effectiveness. Hereby, we assume that people depend on co-workers to be able to do their work (the degree can differ depending on the profession).

Effectiveness guidelines can differ depending on the location. This makes it possible to model e.g. location dependent shop opening hours. They can also combine different effects. For instance, daylight intensity also includes seasonal effects, making it possible to follow e.g. seasonal rhythms.

4 Project Concept

Apart from periodically executed activities (e.g. sleep or daily shopping), people can also have a motivation to execute activities during a certain time period. The motivation and the time period is thereby defined by a special event. An example is the plan to give a party and the necessity to buy extra food before the party starts. In this case, it is the event of having a party that drives people to the shop. We model such events as *projects*. Axhausen (1998) and Schönfelder and Axhausen (2009) define projects as a coordinated set of activities, tied together by a common goal or outcome. Miller (2005) argues that projects are a reasonable organizing principle for dealing with complex human behavior. The presented framework models projects through a mechanism that temporally modifies reference values of behavioral guidelines.

Figure 2: Recursive structure of projects enables composition of bigger and more complex projects

(a) Combination of subproject *organize party food* and tasks *clean apartment* and *prepare food* build project *organize home party*

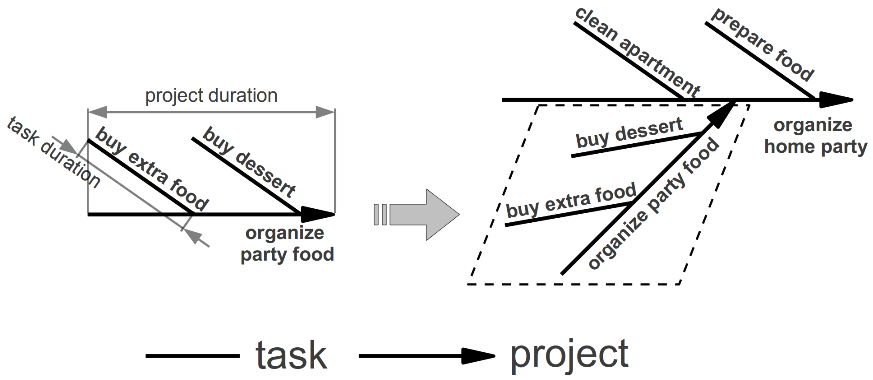

4.1 Structure of Projects

Tasks are the basic components of projects and are linked to activities (e.g. daily shopping). They modify reference values of behavioral guidelines during a specific time period. Table 1 provides several example of possible tasks and their parameters.

A combination of several tasks build a project (see Fig. 2). For instance could the tasks *buy dessert* and *buy extra food* build the project *organize party food*. Projects can also have a recursive structure and contain other projects. The project *organize party food* could be re-used, e.g. for the project *organize home party* which also includes tasks like *prepare food* and *clean apartment*. The project *organize home party* could then become a subproject of *organize wedding* together with other tasks (e.g. *pick up guests*) and other projects (e.g. *organize ceremony*). This concept provides a mechanism where tasks and projects can be re-used to build bigger and more complex projects.

4.2 Reference Value Modification

Projects modify reference values of behavioral guidelines during a specific time period. These changes (difference to previous reference values) are also discounted as it is done for monitoring values (see section Behavioral Model). This is necessary because abrupt changes would cause a sudden increase between reference and monitoring values. Consequently, discomfort would also instantaneously increase, leaving agents with no time to react.

Table 1: Task examples and their parameters

Task	**Parameters**	**Description**
buy dessert	Activity: daily shopping Duration: $+0.25h \pm 0.1$ Frequency: $+1\times^{+0}_{-0}$ Location: confectionery Time: 30/04/2011 from 9:00 am - 4:00 pm	This person needs to do one extra trip (parameter frequency) to the confectionery (parameter location) of approximately 0.25 hours (parameter duration). The task should be done on Sat 30$^{\text{th}}$ Apr, 2011 between 9 am and 4 pm.
buy extra food	Activity: daily shopping Duration: $+0.75h \pm 0.25$ Frequency: $+0\times^{+1}_{-0}$ Time: 30/04/2011 from 9:00 am - 4:00 pm	This person needs to do extra *daily shopping* of approximately 0.75 hours (parameter duration) and is free to combine it with other shopping duties or to do an extra shopping trip (parameter frequency). The task should be done on Sat 30$^{\text{th}}$ Apr, 2011 between 9 am and 4 pm.
get a haircut	Activity: personal care Frequency: $+1\times^{+0}_{-0}$ Location: hair dresser Time: 30/04/2011 from 2:30 pm - 3:30 pm	This person has a hair dresser appointment on Sat 30$^{\text{th}}$ Apr, 2011 from 2:30 pm to 3:30 pm.
pick up guests	Activity: pick up Frequency: $+1\times^{+0}_{-0}$ Location: train station Time: 30/04/2011 at 4:30 pm	This person needs to do one trip (parameter frequency) to the train station to *pick up* guests. The task should be done on Sat 30$^{\text{th}}$ Apr, 2011 at 4:30 pm.
work Saturday morning	Activity: work Duration: $+4.0h \pm 1.0$ Frequency: $+0\times^{+1}_{-0}$ Time: 30/04/2011 from 7:00 pm - 12:00 am	This person needs to work approximately 4 hours on Sat 30$^{\text{th}}$ Apr, 2011 between 7:00 pm to 12:00 am.
work Saturday morning	Activity: work Frequency: $+1\times^{+0}_{-0}$ Time: 30/04/2011 from 9:00 pm - 11:00 am	This person needs to work on Sat 30$^{\text{th}}$ Apr, 2011 from 9:00 pm to 11:00 am. In combination with the work task above, this models flexible working hours with a period when the person must be present at the work place.

5 Planning Heuristic

Other approaches to agent-based microsimulations revealed disadvantages like poor performance for large scenarios (Charypar and Nagel (2006)), high computational costs (Balmer (2007)) or inflexibilities when agents should spontaneously react to unexpected events (Kuhnimhof and Gringmuth (2009)). We consider a planning heuristic as a feasible approach that can overcome these limitations. A heuristic aims to quickly approximate a good solution. Thus, it is unnecessary to have completely accurate knowledge about the state of mind and plans of other agents. This is helpful since we plan to simulate our agents in a distributed computation environment (Charypar *et al.* (2010)) where global knowledge induces extremely high computational costs. A heuristic also enables agents to react to unexpected events because it enables agents to make their decisions spontaneously. One could argue that people seek optimal day plans. However, other authors (e.g. Simon (1955) and Schlich (2004)) doubt that behavior can be explained as a utility maximization function. The aim of this work is to demonstrate how far a decision procedure, that approximate a good solution with limited information, can produce real world behavior.

The next section introduces mathematical formulations used by the planning heuristic during its decision procedure. The section thereafter demonstrates the actual decision steps and the application of the mathematical formulations.

5.1 Mathematical Formulations

5.1.1 Discomfort

Discomfort levels identify the urgency an agent experiences to take action against them. The discomfort an agent receives from an activity at time t is defined as

$$D(t) = \sum_{k=1}^{n} (f^k_{refVal}(t) - f^k_{monVal}(t))^2 \cdot \begin{cases} w^k_1 & \text{if } f_{monVal}(t)_k \leq f_{refVal}(t)_k \\ w^k_2 & \text{otherwise} \end{cases} \tag{1}$$

$$w^k_1 = \frac{1}{(f^k_{refVal}(t) - f^k_{lower-bandwidth}(t))^2} \tag{2}$$

$$w^k_2 = \frac{1}{(f^k_{refVal}(t) - f^k_{upper-bandwidth}(t))^2} \tag{3}$$

the sum of the squared difference of the reference $f_{refVal}(t)$ and monitoring values $f_{monVal}(t)$ of all observation window guidelines n normalized by the squared difference of the reference value and the lower bandwidth w_1^k if $f_{monVal}(t)_k \leq f_{refVal}(t)_k$ or by the squared difference of the reference value and the upper bandwidth w_2^k otherwise.

5.1.2 Discomfort Reduction

The discomfort reduction an agent receives for executing an activity is defined as

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

the difference of the discomfort $D(t_{es})$ at execution start t_{es} and the expected discomfort $D(t_{ee})$ at execution end t_{ee}.

5.1.3 Execution Time Quota

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

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

the integral of the effectiveness guideline $f_{effect}(t)$ between execution start t_{es} and execution end t_{ee} normalized by the execution duration $t_{ee} - t_{es}$. This parameter is important if an agent decides to execute an activity during a time period where it cannot or can only be partially executed (e.g. because the shop closes).

5.1.4 Look-Ahead Measure

Effectiveness guidelines provide information about future execution options. For instance, shop opening hours inform agents about either if they can shop and for how long or how long it takes until they can shop next time. Agents can use such information to plan ahead and e.g. postpone execution of activities because time windows of other activities are going to close soon. Other authors (e.g. Atkinson (1994) and Ioannou *et al.* (2001)) working on *Vehicle Scheduling Problems with Time Window Constraints* also recognized the importance of such information.

The aim of the proposed look-ahead measure is to provide agents with an awareness of decreasing execution options in the near future. We extract this measure through a convolution of the effectiveness guideline with an exponential kernel (see Fig. 3), similar to the convolution of monitoring values (see Fig. 1). In this convolution, we use an exponential kernel that give opening hours in the near future more weight. The look-ahead measure an agent receives for executing an activity at a specific location at time t is defined as

$$LA(t) = \begin{cases} 1 + w_1 \cdot (1 - \int_0^h (f_{effect}(t+x) \cdot kernel(x))\,\mathrm{d}x) & \text{if } f_{effect}(t) > u \\ 1 & \text{otherwise} \end{cases} \tag{6}$$

1 plus the multiplication of w_1 with the difference of 1 minus the integral of the effectiveness guideline $f_{effect}(t)$ multiplied by $kernel(x)$ between 0 and the observation window horizon h if the effectiveness guideline $f_{effect}(t)$ yields a higher value than a predefined threshold u (e.g. 0 for closed shops) or 1 otherwise. Since we use $LA(t)$ as a factor in the final heuristic (see Section 5.1.5), we designed it in such a way that it yields a value in the range of $[1..w_1]$ (1 if execution is not advisable (e.g. because shops are closed) and a value approaching w_1 for decreasing execution options).

5.1.5 Discomfort Reduction Density

Discomfort reduction density is defined as

$$DRD(t_{ts}, t_{es}, t_{ee}) = q(t_{es}, t_{ee}) \cdot \frac{DR(t_{es}, t_{ee})}{t_{ee} - t_{ts}} \cdot LA(t_{ee}) \tag{7}$$

the multiplication of the execution time quota $q(t_{es}, t_{ee})$ with the discomfort reduction $DR(t_{es}, t_{ee})$ between execution start t_{es} and execution end t_{ee} normalized by the difference between execution end t_{ee} and travel start t_{ts} multiplied by the look-ahead measure $LA(t_{ee})$ at execution end. Including execution time quota $q(t_{es}, t_{ee})$ ensures that activities which cannot or can only be partially executed (e.g. because the shop closes) get a lower discomfort reduction density. Normalizing with the travel duration ensures that locations which are further away from the current location or need a long time to reach because of traffic congestions receive a smaller discomfort reduction density. This provides for a simplistic location choice procedure with agents preferring locations close to their current location. Including the look-ahead measure $LA(t_{ee})$ ensures that activities with fewer execution options in the future receive a higher discomfort reduction density and are therefore preferred by agents.

5.2 Decision Procedure

The planning heuristic uses a two-step decision procedure to determine the activity-location pair an agent should execute next. In a first step, it identifies promising activity-location pair

Figure 3: Illustration of look-ahead measure for shop opening hours with short and long weekends.

(a) Look-ahead measure with a kernel of 2 days. The higher the measure the closer the end of the current shop opening window. The measure is higher before weekends indicating less shopping options in the near future.

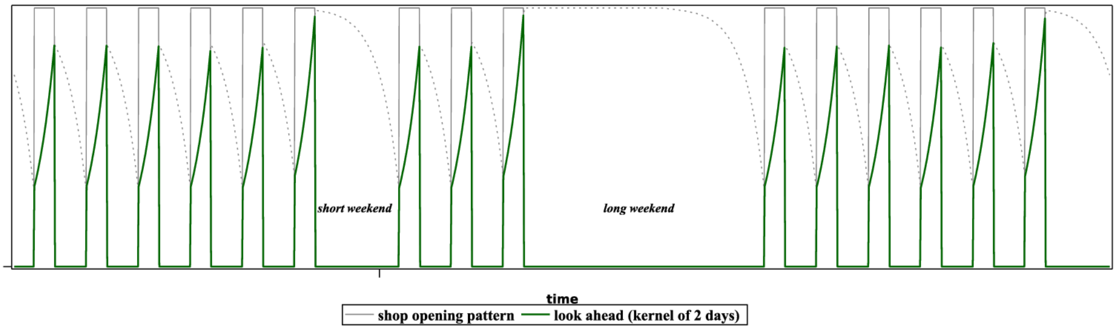

(b) Look-ahead measure with a kernel of 7 days. This kernel can differentiate between short and long weekends (measure is higher before long weekend). Choosing the right kernel length is important and we propose a duration of approximately 2 to 3 times the average interval between two activity executions (e.g. $3 \cdot 2$ days = 6 days for daily shopping).

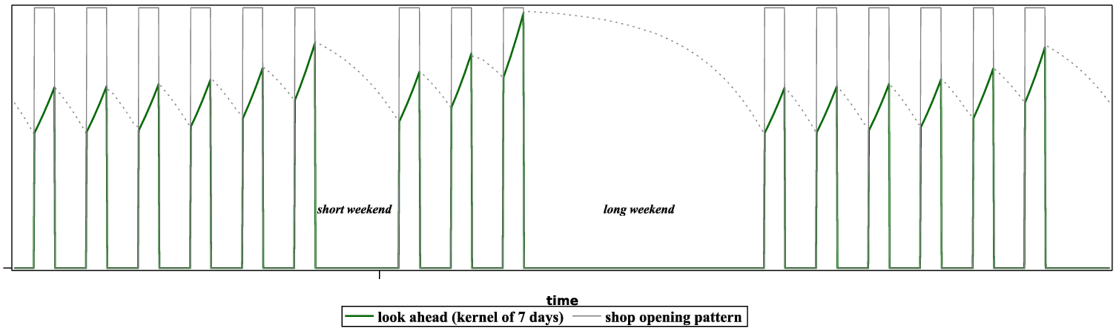

candidates. Here, the planning heuristic makes *best guesses* for values which are expensive to compute. In a second step, the planning heuristic computes optimal values of promising candidates and decides to implement the most promising activity-location pair.

5.2.1 First Step

In the first step, the planning heuristic makes the following assumptions to determine promising activity-location pair candidates:

- It uses the free speed travel time to compute travel durations between locations. Computing the exact travel time is expensive because it depends on the current time and could include different computer nodes (we plan to run the simulation on a distributed computation environment (Charypar *et al.* (2010))). We reduce deviations between exact travel time and free speed travel time through a multiplication with a factor. The current version of the simulation uses a constant factor (e.g. 1.2) but we consider a learning process for later versions where agents adapt the factor based on their past experience (e.g. by time-of-day, type-of-location etc.).

- It uses the reference activity duration of the execution duration guideline (see Section 3.1.2) as the planned duration. Determining the optimal duration requires a numerical optimization which is computationally expensive.

The planning heuristic computes the discomfort reduction density $DRD(t_{ts}, t_{es}, t_{ee})$ for all activity-location pairs using the above mentioned assumptions.

5.2.2 Second Step

In the second step, the planning heuristic uses real travel durations and computes optimal execution durations (using the discomfort reduction density function $DRD(t_{ts}, t_{es}, t_{ee})$) for the 20% most promising activity-location pair candidates of the first step. The optimization of the execution duration is done numerically using Brent's method from Press *et al.* (2007).

Since the optimization is computationally expensive, the heuristic tries to reduce the range of valid execution durations before it starts the optimization. The ratio of the guidelines *percentage of time* and *frequency* defines the average execution duration (see Fig. 4). When an agent decides to execute an activity, this ratio instantaneously drops and recovers during activity execution. The time to recover into the bandwidth of the duration guideline (see Section 3.1.2) yields lower and upper duration bounds.

Effectiveness guidelines (see Section 3.1.3) can further narrow valid upper duration bounds.

$$duration{-}bound_{upper} = \begin{cases} f_{effect}(y)^{-1} & \text{if } f_{effect}(y)^{-1} < duration - bound_{upper} \\ duration - bound_{upper} & \text{otherwise} \end{cases} \tag{8}$$

The upper duration bound is updated with the time $f_{effect}(y)^{-1}$ when the effectiveness guideline drops below a predefined value y (e.g. to 0 because shops close), if this time is earlier than the current upper duration bound.

Finally, the planning heuristic searches the optimal execution duration within the lower and upper duration bound and implements the activity-location pair that yields the highest discomfort improvement density.

Figure 4: Illustration of average execution duration

(a) The average execution duration instantaneously drops when an agent decides to execute an activity. The time to recover into the bandwidth of the duration guideline specifies lower and upper duration bounds

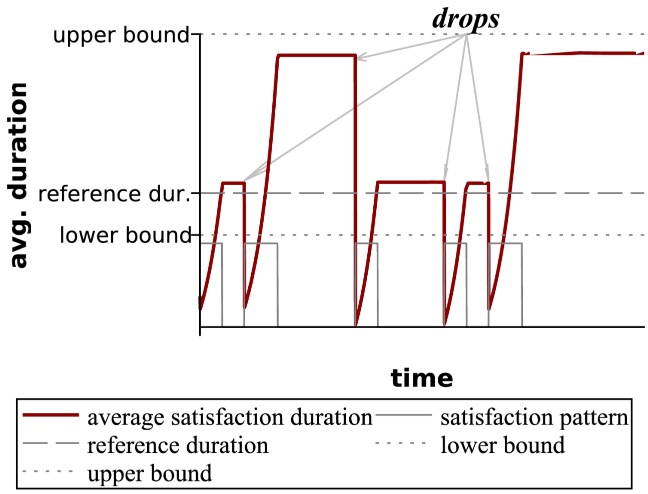

6 Outlook

The next task is to finalize the implementation of the proposed concepts. We program in C++ since performance is important and we can build on existing code from Märki *et al.* (forthcoming). In the first validation phase, we will focus on the heuristic and calibrate its parameters to see if it is able to produce realistic behavior. This includes reoccurring activities and activities modeled through projects. Results will show the necessity of a heuristic adaptation. In a second phase, we will validate our model by comparing simulation data with analyses of two existing six-week continuous travel diaries (Axhausen *et al.* (2002), Löchl *et al.* (2005) and Schönfelder (2006)).

7 Conclusion

This paper proposes a microscopic travel demand simulation that can continuously simulate agent's behavior and resulting movements. Behavioral guidelines are central for the proposed behavioral model. These guidelines are closely related to statistical data provided by various sources (e.g. Swiss Federal Statistical Office (2006)), simplifying model utilization for practitioners. We illustrate different guidelines and their parameters. Some guidelines have observation windows enabling performance tracking over different time horizons. Time-dependent effectiveness guidelines model various effects like shop opening hours or social and cultural norms. We propose using projects to model non-recurring tasks. Projects are limited to a specific time period during which they influence behavioral guidelines. Agents keep track of their performance and compare it to behavioral guidelines. Deviations cause discomfort which is conveyed to a planning heuristic, making on the fly decisions about upcoming activities agents should execute.

8 References

Arentze, T. A. and H. J. P. Timmermans (2006) A new theory of dynamic activity generation, paper presented at the *85th Annual Meeting of the Transportation Research Board*, Washington, D.C., January 2006.

Arentze, T. A. and H. J. P. Timmermans (2009) A need-based model of multi-day, multi-person activity generation, *Transportation Research Part B: Methodological*, **43** (2) 251–265.

Atkinson, J. B. (1994) A greedy look-ahead heuristic for combinatorial optimization: An application to vehicle scheduling with time windows, *Journal of the Operational Research Society*, **45** (6) 673–684.

Axhausen, K. W. (1998) Can we ever obtain the data we would like to have?, in T. Gärling, T. Laitila and K. Westin (eds.) *Theoretical Foundations of Travel Choice Modeling*, 305–323, Pergamon, Oxford.

Axhausen, K. W., A. Zimmermann, S. Schönfelder, G. Rindsfüser and T. Haupt (2002) Observing the rhythms of daily life: A six-week travel diary, *Transportation*, **29** (2) 95–124.

Balmer, M. (2007) Travel demand modeling for multi-agent traffic simulations: Algorithms and systems, Ph.D. Thesis, ETH Zurich, Zurich, May 2007.

Charypar, D., A. Horni and K. W. Axhausen (2009) Need-based activity planning in an agent-based environment, paper presented at the *12th International Conference on Travel Behaviour Research (IATBR)*, Jaipur, December 2009.

Charypar, D., A. Horni and K. W. Axhausen (2010) Pushing the limits: A concept of a parallel microsimulation framework, *Working Paper*, **640**, IVT, ETH Zurich, Zurich.

Charypar, D. and K. Nagel (2006) Q-learning for flexible learning of daily activity plans, *Transportation Research Record*, **1935**, 163–169.

Gliebe, J. P. and K. Kim (2010) Time-dependent utility in activity and travel choice behavior, *Transportation Research Record*, **2156**, 9–16.

Ioannou, G., M. Kritikos and G. Prastacos (2001) A greedy look-ahead heuristic for the vehicle routing problem with time windows, *Journal of the Operational Research Society*, **52** (5) 523–537.

Kuhnimhof, T. and C. Gringmuth (2009) Multiday multiagent model of travel behavior with activity scheduling, *Transportation Research Record*, **2134**, 178–185.

Löchl, M., S. Schönfelder, R. Schlich, T. Buhl, P. Widmer and K. W. Axhausen (2005) *Untersuchung der Stabilität des Verkehrsverhaltens*, 1120, Eidgenössisches Departement für Umwelt, Verkehr, Energie und Kommunikation, Berne.

Miller, E. J. (2005) An integrated framework for modelling short- and long-run household decision-making, in H. J. P. Timmermans (ed.) *Progress in Activity-Based Analysis*, 175–201, Elsevier, Oxford.

Märki, F., D. Charypar and K. W. Axhausen (forthcoming) Continuous activity planning for a continuous traffic simulation, *Transportation Research Record*.

Nagel, K. and G. Flötteröd (2009) 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.

Press, W. H., S. A. Teukolsky, W. T. Vetterling and B. P. Flannery (2007) *Numerical Recipes: The Art of Scientific Computing*, 3. edn., Cambridge University Press, Cambridge.

Schlich, R. (2004) Verhaltenshomogene Gruppen in Längsschnitterhebungen, Ph.D. Thesis, ETH Zurich, Zurich, April 2004.

Schönfelder, S. (2006) Urban rhythms: Modelling the rhythms of individual travel behaviour, Ph.D. Thesis, ETH Zurich, Zurich.

Schönfelder, S. and K. W. Axhausen (2009) Travel as a function of (life) projects, paper presented at the *European Transport Conference*, Leeuwenhorst, October 2009.

Simon, H. (1955) A behavioral model of rational choice, *Quarterly Journal of Economics*, **69**, 99–118.

Swiss Federal Statistical Office (2006) *Ergebnisse des Mikrozensus 2005 zum Verkehrsverhalten*, Swiss Federal Statistical Office, Neuchatel.