

A GA-based household scheduler

KONRAD MEISTER*, MARTIN FRICK & KAY W. AXHAUSEN

*Institute for Transport Planning and Systems (IVT), Swiss Federal Institute of Technology (ETH)
Zurich, CH – 8093, Zurich, Switzerland*

*(*Author for correspondence, E-mail: meister@ivt.baug.ethz.ch)*

Key words: activity scheduling, agent-based demand generation, genetic algorithms, intra-household interaction

Abstract. One way of making activity-based travel analysis operational for transport planning is multi-agent micro-simulation. Modelling activity and trip generation based on individual and social characteristics are central steps in this method. The model presented here generates complete daily activity schedules based on the structure of a household and its members' activity calendars. The model assumes that the household is another basic decision-making unit for travel demand aside from individual mobility needs. Results of the model are schedules containing complete information about activity type and sequence, locations, and means of transportation, as well as activity start times and durations. The generated schedules are the outcome of a probabilistic optimisation using genetic algorithms. This iterative method improves solutions found in a random search according to the specification of a fitness criterion, which equals utility here. It contains behavioural assumptions about individuals as well as the household level. Individual utility is derived from the number of activities and their respective durations. It is reduced by costs of travelling and penalties for late, respectively early arrival. The household level is represented directly by the utility of joint activities, and indirectly by allocation of activities and means of transportation to household members. The paper presents initial tests with a three-person household, detailing resulting schedules, and discussing run-time experiences. A sensitivity analysis of the joint utility parameter impact is also included.

1. Scheduling households

The study of activities generation and associated trips has often been somewhat neglected in travel behaviour research. This is not surprising; the comparison of static (one-day) cross-sections was the traditional focus of analysis. The two most relevant recent advances were the idea of homogeneous groups of travellers and their associated activity chain distributions (Poeck and Zumkeller 1976; Axhausen 1990; Fellendorf et al. 1997) and the successful integration of the activity pattern choice into random utility models (Bowman 1998; Bowman et al. 1998). The first approach allows the representation of overall activity patterns variability in a simulation, but does not provide for a mechanism to change either the activity patterns themselves, or their distributions

for the particular homogeneous group. The patterns and generalized costs are linked in the model structures of the second approach, but their effects tend to be small and dominated by the effects of the socio-demographic variables, which often reduce the models to sophisticated cross-classifications. In addition, as utility maximizing approaches, one must impose the assumption that the daily pattern is planned in one step at one point in time. It is obvious using introspection and well known from empirical work since the late 1970s, that this assumption is unrealistic, even if computationally and econometrically necessary (Jones et al. 1983). In particular, Doherty's work with his activity scheduling surveys has shown that a large share of activities undertaken is planned during the day or even started at short notice, i.e. spontaneously (Doherty and Miller 2000; Rindsfuser et al. 2003 and references therein).

Two developments have now increased the focus on trip generation again: first, the interest in household interaction and the resulting allocation of tasks to different household members; second, the interest in responses to real time traveller information and traffic management systems, including tolling. Both are critical to activity scheduling, i.e. the determination of timing and duration of an activity and its associated characteristics (location, mode, group size, expenditure), either as a bargaining process within the household or as a response to information or pricing signals. This process of activity scheduling involves the identification of activity needs, instances of activity suppression, and activity scheduling and re-scheduling over arbitrary time horizons and is therefore inherently dynamic.

The concept of scheduling has thus become central to the current work in activity-based analysis (Damn 1983; Jones et al. 1990; Kitamura 1996; Pas and Harvey 1997). Four lines of work can be identified:

- empirical work observing the scheduling process through tracking plans and realized behaviour (see above),
- stated-response work which attempts to understand the scheduling process as it happens (Ettema et al. 1997),
- econometric estimates of utility functions and other attempts to reproduce and predict activity patterns (Recker 1999; Arentze and Timmermans 2000; Glibe and Koppelman 2002; Joh 2004; Zhang et al. 2004)
- conceptual and simulation models of the activity scheduling process (e.g. Gärling et al. 1989; Axhausen and Goodwin 1991; Gärling et al. 1994; Doherty and Axhausen 1998; Gärling et al. 1998; Charypar and Nagel 2003; Roorda and Miller 2004).

In addition, the work on departure time choice is related to this material through implied trade-offs between the two activities at the start and the end of the trip (see Noland and Small 1995 and references therein).

This paper will contribute to the fourth stream of work by presenting a scheduler based on genetic algorithms that can accommodate an arbitrary number of household members. Processes modelled at the household level are: division of work, joint activity participation, and allocation of means of transportation to household members. Considering these factors, the scheduler creates individual daily activity schedules. This study attempts to make the point that social networks (in general) and households (in particular) have a significant influence on travel activity scheduling.

The paper is structured as follows. Before describing the scheduler and proposed solution for household interactions, the next section places the model into a dynamic framework to sketch the current limitations of the approach. After the description of the model and its capabilities, the paper will discuss the results for an example household with three members. An extra-section is devoted to the sensitivity analysis of the parameter rewarding additional utility of joint activities. The final section discusses run-time experiences and future research needs.

2. A dynamic framework

The individual's position in space at any one time is the product of his or her biography. In a first step, it is useful to concentrate on activity scheduling for normal, daily tasks. We assume that longer-term choices, such as partnership, children, home and work locations, but also the available set of mobility tools (motorized vehicles, bicycles, public transport season and discount tickets) are fixed in the short term. The elements forming one's personal knowledge of the world can be approached in two ways, each stressing a different facet of that knowledge: On one hand, the *mental map* describes locations, their relative positions, and networks and routes linking them with associated generalized costs reflecting the person's experiences and expectations (Lynch 1960; Stern 1990). On the other hand, the *activity repertoire* consists of locations and types of activity which can be performed there for given generalized costs and which return a particular level of satisfaction. It is worth pointing out that the number of locations included in the mental map/repertoire is substantially larger than the observed set of locations for a given time period, which constitute the activity space (see Schönfelder and Axhausen 2003). A further element of the personal world is the set of expectations that the traveller has formed about the world based on his/her experiences. These assumptions determine his/her assessment of parts of the network, which he/she has not visited so far. Parts of these expectations are search rules and methods, which are available to answer questions about daily activities, their locations, costs, and expected satisfaction levels.

The performance of everyday life is limited by the personal world, which is continuously evolving (see Figure 1). Physiological needs, higher level needs, commitments accepted, but also the seasons activate elements in the activity repertoire, which in turn could satisfy these needs or allow the traveller to fulfill his/her obligations. These activated elements constitute the activity calendar, which evolves in response to needs, desires, and activity execution, but also in response to unforeseen opportunities. The entries in the calendar have priorities, which respond to the level of need or commitment to the activity (type). As a rule, there is no one-to-one match between an activity type and need, or even commitment. The substitutability should be lower for commitments, such as work or the care of a person, as commitments are more specific in their description of the necessary activities than the need for relaxation, which can be achieved in many ways. Entries in the 'things-to-do'-calendar, in general, will involve more than can be accommodated during the next day or generally, short-term planning period. It is important to note that the traveller's commitments and projects provide activities for some time into the future (Axhausen 1996). The activity calendar will therefore contain certain activities that need not to be started until well after the next day. Survey work is now only started to provide empirical evidence on the frequency, duration, and effort involved in the various projects that persons and households undertake at any one time. A project is a set of activities linked through a common goal. It can be trivial, such as preparing dinner, or non-trivial, such as buying a new house and moving. The concept of the project is important, as it links various disparate

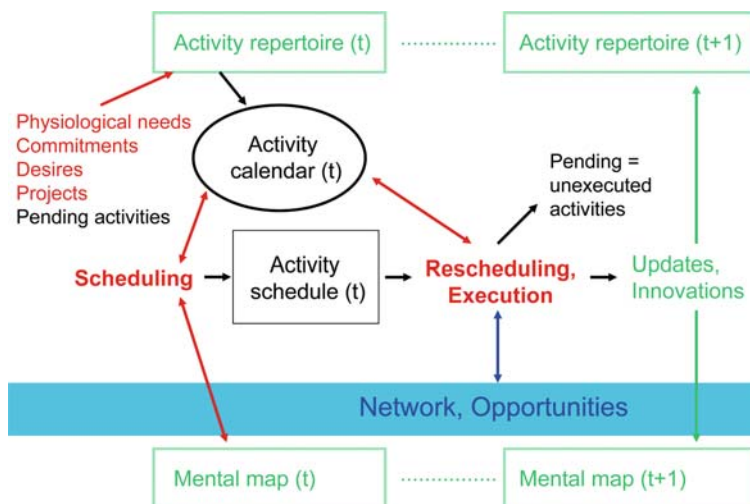


Figure 1. Daily scheduling within a dynamic framework. Source: (Kitamura R 1996).

activities into a coherent whole and provides the time frames for subsidiary activities.

Through organizing and planning, travellers create schedules to guide a given day. Scheduling must provide time for high priority activities in the calendar, fill any remaining gaps, and satisfy commitments on the basis of a person's knowledge of his environment or what he/she can find out about it in reasonable time. It is the central management process of everyday life and its dynamics.

The conceptual framework of Figure 1 sees scheduling taking place at the start of the day, but continuing throughout the day, when parts of the schedule have been addressed. Changes might be necessary because the transport system has been unreliable, an event was cancelled, a store did not have the desired item, etc.

At the end of the day/planning period, the traveller is left with unexecuted activities, which he/she needs to carry over into the next day, cancel (suppress), or reshape. In addition, the traveller has updated the mental map against current experiences, and incorporates innovations, which he/she might have developed during the day or participated in on the suggestion of a third party: a new restaurant, new public transport connection, new parking garage, etc.

This conceptualization does not specify what form scheduling takes. Rather, it assumes that scheduling is not, strictly speaking, optimal, but that it is heuristic and preliminary, accounting for the 'expected unexpected'. In line with Doherty's and others' results, one would assume that the scheduling traveller leaves empty spaces in his plans and cannot achieve the equilibrium conditions of time-use economics (Becker 1965; Deserpa 1971; Jara-Diaz and Farah 1988; Bates et al. 1996).

3. Model structure

Most schedulers implemented in software so far have, in spite of the apparent flaws in this assumption, enforced the equilibrium condition on the marginal utilities of activity participation (see Charypar and Nagel 2003; Joh 2004). The computational advantages of the assumption argue for this. Random utility choice models, which impose the same assumption in the estimation of the parameters of utility functions, suggest it as well. The model presented here follows this tradition. According to the conceptual scheme of Figure 1, the model has to select activities from the calendar and schedule them completely (timing, duration, mode, location and joint participation of other household members). It is consistent with the concept, when not every activity present in the calendar is scheduled for the day in question.

The model presented here is an extension of an existing genetic algorithm (GA) scheduler developed by D. Charypar at ETH (Charypar and Nagel 2003). This further development attempts to overcome the limitation of Charypar's model to separate individuals by allowing household members' schedules to interact and to synchronize to some degree. At first, we assume that household members like to do things together. The basic instrument is to extend the individual utility function to account for additional utility derived from joint participation in certain activities. Furthermore, activities can be allocated to household members to represent division of work. A third intra-household interaction concerns transport: availability of transportation alternatives for household members, specifically their varying users during the day. Some new dimensions of individual choice are added to the scheduler. The model will be referred to as 'the household scheduler'.

To deal with the complexity of the resulting household utility function, the household scheduler employs the GA approach to search for good solutions (Goldberg 1989). The GA's basic objective is to iteratively improve an initially random set of individuals, in our case household schedules. The instruments for this search are the cross-breeding of selected good schedules (cross-over) and their further, slow, random mutation. In a constant-size population of schedules, only those individuals with high fitness are selected, while bad schedules are dropped. Experience has shown that a GA is able to find near optimal solutions after a sufficient number of iterative applications of these two instruments, even for very badly behaved utility surfaces.

The model has to specify the coding of the schedules, describe the elements (activities) of the calendar, and define a suitable utility function. Please refer to the study of the previous model (Charypar and Nagel 2003) for comparison as it served as a starting point for the work presented here.

3.1. *Model input*

The main input variable to the household scheduler is the household members' activity calendar. Table 1 shows two example activities and their dimensions. Most of the variables are used in the utility function described later.

The reference to a basic need is a new element here. Certain activities have clear rhythms over time, such as sleeping or eating. The utility function cannot assure this spacing in its naïve form. Therefore, the scheduler discounts the utility of such an activity with a logit curve, if it occurs too quickly after a previous activity satisfying the same need. Over longer time horizons, these rhythms are also evident for other activity types Schönfelder

Table 1. Activities and their dimensions in the activity calendar.

Variable name/unit	Description	Example 1	Example 2
Purpose	Activity (mandatory)	Breakfast	Soccer
Type	Activity type: Individual, allocated (to one person) or joint activity (mandatory)	Joint	Individual
Participants	Possible participants in the activity specified as a list of the relevant household members A, B, C,...(mandatory for joint and allocated activities)	ABC	–
Priority P_i []	The higher the priority, the higher the utility level of the activity, and the higher its probability to be scheduled. Values are 1, 2, 3,...with 1 being the highest priority (mandatory)	2	3
$t_{w, i}$ [h]	Working point of activity i . All activities have the same marginal utility at their working points. The closer the durations of all scheduled activities are to their working points, the better the schedule (mandatory)	1	2
$t_{late, ar, i}$ [time]	Latest arrival (start) time for the activity (optional)	8:00	–
$t_{early, dp, i}$ [time]	Earliest departure (end) time of the activity (optional)	10:00	–
$t_{short, i}$ [h]	Minimum duration (optional)	0.75	1
$\beta_{joint, i}$ [Arentze and Timmermans 2000]	Parameter of the utility of joint activity performance (mandatory for joint activities)	0.2	–
Need	Association of the activity with a basic need (optional)	Hunger	–
Facility	Facility type suitable for the activity. There are many locations for one facility. It is possible to specify if an activity type can only be undertaken at a particular location, such as home or a pre-allocated work place or kindergarten (mandatory)	Home	Leisure

and Axhausen 2001; Bhat et al. 2004 The household activity pattern problem: General et al. 1999). In this stage of development, only hunger is implemented in the model.

The variable *Participants* allows the modeller to describe allocations of an activity to persons, or their participation in joint activities. For allocated activities, only one of the possible participants may perform a scheduled activity. For a possible joint activity, the modeller specifies which household members may participate. Both types of activities are listed in the household part of the calendar, while each household member has his/her own calendar with individual activities.

The household description also contains a list of means of transportation available. The scheduler does not incorporate mode choice as a separate model. The idea is to impose constraints on household members' scheduling opportunities by defining cars or bikes as 'tied vehicles'. If one household member plans to use a tied vehicle, it will not be available to others. The dynamics of mode choice within the household are determined by their generalized costs, their average speeds, and availability. Every mode tied to a vehicle has an additional constant to generally represent access time. This approach reflects the fact that decisions about possession and usage of mobility tools are in general made at the household level. Since mode choice is an element of the activity encoding, it is part of the optimization with the GA (described later).

A further critical input to the scheduler is a map describing where the agents' daily lives take place. As shown in Figure 4, the environment is loosely modelled on Karlsruhe, as found in the Mobidrive survey (Axhausen et al. 2002), which is intended to serve as the validation data source in later stages of this work. The size of the environment determines the size of the choice set for location choice of secondary activities: Each agent has on the one hand, fixed locations for long-term commitments, so-called primary activities, including home, work, school, etc. On the other hand, secondary activities may be performed in different locations of a certain facility type. The setup described in the Results section has 30 locations for each shops and leisure facilities. There is no utility connected to them, so their choice still happens completely randomly. On the presented map, 160 households could be placed and scheduled. All distances are calculated crow-fly, since the scheduler is not yet integrated into a transportation network model.

3.2. *Utility function*

The fitness F to be maximized by the GA is defined as the household utility function HUF. It is the unweighted sum of the household members' individual utilities:

$$F = \text{HUF} = \sum_m U_{\text{total},m}$$

There is little empirical guidance for the formulation of the utility function

of an activity (see Axhausen 1990 or Joh 2004 for reviews). Reflecting current assumptions, the model includes the following elements:

- Positive utility $U_{dur,i}$, derived from performing an activity i .
- Travel, late arrival, early departure and the violation of the minimum time constraint are penalized linearly as function of their respective durations t_{travel} , $t_{late.ar}$, $t_{early.dp}$ and t_{short} .
- Opportunity costs arise by waiting periods, e.g. for a shop or a leisure facility to open.

See Table 2 for values of the respective parameters. So the total utility of household member m planning to perform k activities is

$$U_{total,m} = \sum_{i=1}^k U_{dur,i} - (\beta_{travel} t_{travel} + c_{travel}) - \beta_{late.ar} t_{late.ar,i} - \beta_{early.dp} t_{early.dp,i} - \beta_{short} t_{short,i}$$

The contribution of the performance of activity is logarithmic to the ratio of its actual duration t_{dur} and the break-even duration t_0 , which is scaled to reflect its priority P_i :

$$U_{dur,i} = \beta_{dur,i} t_w \ln\left(\frac{t_{dur,i}}{t_0}\right)$$

$$t_0 = t_w e^{\left(\frac{-c_i}{\beta_i t_w}\right)}$$

Table 2. Utility function parameters.

Utility element	Symbol	Value	Effective disutility ^a
Activity performance	β_{dur}	20 €/h ^b	–
Waiting	β_{wait}	0 €/h	–20 €/h
Travel costs	β_{travel}	20 €/h	–40 €/h
Late arrival	$\beta_{late.ar}$	60 €/h	–60 €/h
Early departure	$\beta_{early.dp}$	20 €/h	–20 €/h
Violation of minimum duration constraint	β_{short}	$\beta_{early.dp}$	–20 €/h
Joint performance	β_{joint}	0.1	–
Fixed access costs for tied vehicles	c_{travel}	Car: 1.50 € Bike: 0.15 €	Car: 1.50 € Bike: 0.15 €

Source: Similar to the utility concept described in [Schönfelder and Axhausen 2003].

^a At the working point, travelling and waiting create additional opportunity costs of $\beta_{dur} = 20$ €/h. Late arrival, early departure and the violation of minimum duration constraints are penalized directly, therefore not discounted with β_{dur} .

^b This is the approximate average Swiss wage after taxes and health insurance.

This basic formulation is extended in two ways. At first, joint performance of joint-type activities increases the utility. The more a person m is synchronized with other household members in n such an activity, that is the closer $s_i^{m,n}$ is to 1, the higher the additional utility:

$$U_{dur,i} = U_{dur,i} \cdot (1 + \beta_{joint,i} \sum_{\forall n \neq m} S_i^{m,n}); \quad S_i^{m,n} \in [0; 1]$$

Given current knowledge about the effects of such joint performance, we have assumed a simple linear impact. Other forms (log, logit, etc.) are possible, even likely (Zhang et al. 1971). The second modification reflects the urgency for performing an activity derived from the current level $n_{i(t)}$ of its associated need:

$$U_{dur,i} = U_{dur,i} \cdot n_i(t); \quad n_i(t) \in [0; 1]$$

The more activities the algorithm is able to schedule, the higher the fitness will be. This must not be misunderstood as “The more packed my day, the better”, since sleep or relaxation are possible activities, too. There is no penalty for not including all activities present in the calendar. Allocated type activities do not explicitly occur in the utility equations. The algorithm simply checks that not more than one household member performs such an activity. Its utility is calculated as described.

4. GA details

4.1. Encoding

For the GA, each household member’s activities are encoded with the following five elements:

- *Scheduled (binary)*: This information defines whether an activity will be a part of the schedule or not. In the latter case, it is ignored during utility calculations.
- *Sequence (integer)*: Here, the order of activities in the schedule is encoded. Purpose is known from the description in the calendar.
- *Location (integer)*: There are potential multiple locations for each facility. For daily activities’ location choice, a new location is chosen for every new schedule. Primary activities have fixed locations.
- *Time allocation (real double)*: What share of a 24-hour-day is reserved for this activity? The algorithm divides the allocated time for the activity between travel time, possible waiting periods, and performance of the activity. For this, it considers distance, chosen mode of transport, and environmental constraints, e.g. opening times.

- *Mode (integer)*: For each activity, a mode of transportation is chosen. If a trip is necessary to that activity, the mode choice is considered in utility calculation.

For a two-person-household, a complete schedule with eight activities each on the calendar is thus a vector of up to (2 persons·8 activities/person·5 variables/activity) 80 variables. When creating the initial population, all of these variables are set randomly. Therefore, an initial population will mainly consist of nonsense schedules. During evolution, their values depend on the previous schedules, except for mutations taking place during the iterations.

The population of schedules was classified by three of the five variables (sequence, location and mode for all household members and scheduled activities). If a second schedule of the same class is generated, the one with higher fitness is kept in the population, while the other one is deleted. If a second schedule with almost the same fitness is generated, it is also deleted. By ignoring timing information in schedule classification, the population variance is increased. This reduces the ability of the GA to optimize departure times and activity durations, but also minimizes the danger of getting stuck in a local optimum. In this implementation, this trade-off was accepted, as it was felt that it was more important to cover the solution space with a small number (50) of schedules.

4.2. Parameters

Table 3 gives a brief overview of the various GA parameters that have to be configured. All these parameters have to be chosen, according to the nature of the problem to be solved, which is often done on a gut level, so in this case. A necessary step to be done is to determine their sensitivity to the results using an exact method, e.g. a Monte Carlo procedure or another GA.

5. Results

5.1. Example scenario

In this development phase, the scheduler was applied to a number of household types and related activity calendars. As an example, the schedule of a three-person household on a working day is presented here. The respective activity calendars reflect their socio-demographic status:

Table 3. GA parameters.

Parameter name	Description	Value
Popsize	Constant population size.	50
Ngen	Number of generations. It may be used as the maximum number of generated individuals. Here, it serves as a reference point to the stop criterion described below.	100,000
Stop criterion	Some boolean expression. If it becomes true, the GA will stop to evolve the population, and will output the final result.	Here, the GA stops when average fitness didnot increase more than 1% during the last 10,000 iterations.
P_{mut}, τ_{mut}	Probability that one part of a new schedules' encoding will mutate according to its respective mutation operator. Each time a new individual was inserted into the population, P_{mut} is adapted. The higher τ_{mut} , the faster P_{mut} decreases.	Initial: 30%, exponentially decreasing to minimum 7%.
Mindiff	Minimum fitness difference between two individuals. If a new schedule with almost the same fitness is generated, it will be dropped.	0.1

- Joint activities include mealtimes (“breakfast”, “lunch”, “dinner”) and a leisure activity (“chillout/relaxing”). All joint household activities take place at home.
- Shopping is the only allocated activity considered for execution on that day.
- The first person (“academic”) is an academic employed in the city center of the synthetic region. He has to complete a full workday, with the possibility of a leisure activity should be there in the remaining time. Furthermore, he has the opportunity to have his lunch and dinner in the cafeteria which is located right at his workplace.
- His wife (“housewife”) is not employed and stays at home during most of the day. Her individual activity calendar also includes a leisure activity out of home.
- Their child (“young pupil”) who has a high-priority school attendance in the morning. Other possible activities are homework and soccer. The modelled family lives in a small village about 12 km away from the city, and owns one car, which must be shared between the two adult household members.

The example results shown here will be discussed with respect to the performance of the GA. They are visualized in three styles (see Figures 2–5):

- *Time use by person*, as pioneered by the Household Activity Travel Simulator (HATS), the survey tool of Jones (Jones 1979). It shows, in a compact image, where each person is over the course of the day, in a view that aggregates facility types (home, out-of-home, and travel),
- *Time use by location* for all household members as a group. This allows easy detection of joint activities, but also indicates opening times of facilities as environmental constraints on activity scheduling,
- Sequence of locations visited over the course of the day on a *map of the environment* for each person. The legend appended explains the respective elements of the figures.

The three-person example was chosen because it contains each intra-household interaction modelled:

- *Joint activities*: From the potentially three joint mealtimes, only the dinner could be scheduled as a joint activity for all three household members. The housewife and her child share time at home in the afternoon. The academic takes his lunch at work, because it would take too much time to go home for it.
- *Division of work*: The algorithm assigns the shopping task to the housewife who accomplishes it around noon before attending her leisure activity.
- *Allocation of means of transportation*: The respective deterrence function lets the algorithm to assign the family car to the person with longer trips during the day. In this case, the employee who works in the city will use it. It therefore is not available to his wife, who has to walk through the village, where the family lives.

In this scenario, the activity calendars were of different size. For example, the academic has 15 activities on his calendar, with 5 of them being opportunities to have a meal: three at home, two in the cafeteria. Three of those were scheduled. This is enough to survive, and shows the effect of the hunger need associated with a mealtime activity. On the opposite, the other two persons' calendars are much smaller. This accelerates convergence, but decreases behavioural diversity.

The performance of the algorithm is measured with two variables (see Figure 6):

- *Average fitness across entire population*,
- *Standard deviation of schedules' marginal utilities*. This captures how optimal the simulated schedules are. As pointed out, a schedule can be

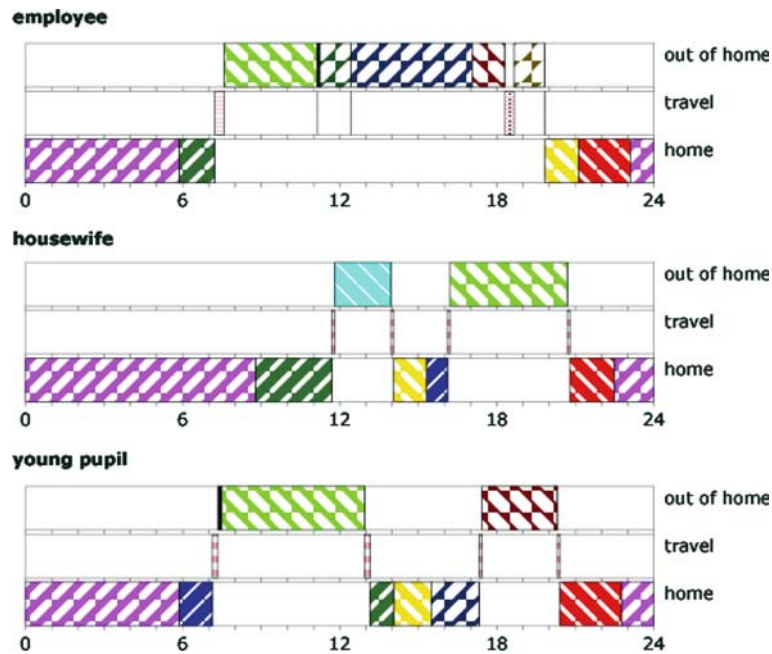


Figure 2. Three-person household – time use by person.

called optimal, when the durations of the performed activities are close to their working points t_w , where all activities have the same marginal utility β_{dur} . Thus, good schedules have two main properties: many activities scheduled (high fitness) and optimal time allocation (low sd of marginal utilities of all performed activities). This measure is used to evaluate the algorithm's convergence. It shows a more detailed picture of what happens inside the population, e.g. when new regions in the solution space are found.

The performance analysis shows the expected behaviour. Average utility is by definition, a monotonically rising numerical series. Its gradient decreases constantly until the stop criterion is reached after 170,000 generations (see GA parameter description). The differentiation described earlier, which leaves out time information when calculating similarity, ensures that variability will not disappear completely. The convergence measured as the standard deviation of marginal utilities shows a quickly oscillating behaviour at the beginning. Schedules that are not completely nonsense are found after a third of the optimization. The optimality levels of the agents' schedules converge each other as expected, but they do not converge to zero. This is

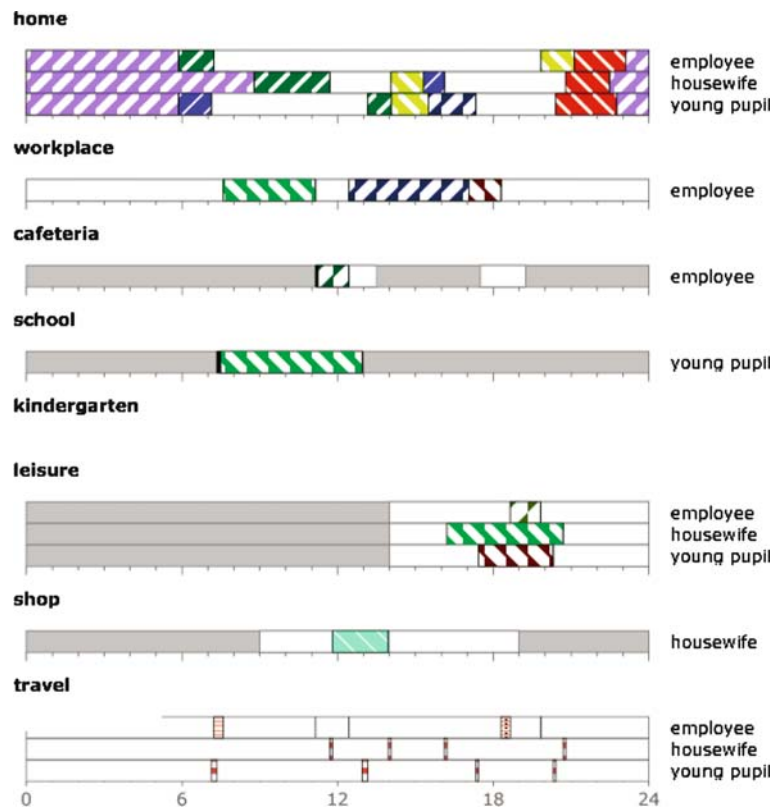


Figure 3. Three-person household – time use by facility.

because of the GA selection criterion described which does not explicitly optimize departure times and activity durations.

5.2. Sensitivity to joint activities' valuation

Although the scheduler employs dozens of variables, we want to concentrate on model response to systematic changes of the parameter rewarding joint activity performance.

The introduction of β_{joint} , and therefore the explicit valuation of joint activity, raises the question whether the schedules respond to its change in the expected way. The expected increase in the joint time and in synchronization (not shown, as it is closely parallel) is observable (Figure 7). The value of $\beta_{\text{joint}} = 0.1$ (as used in the example) more than doubles the joint time in comparison to a zero valuation for joint time. While joint activity durations rise with a higher value of β_{joint} , fitness starts to decrease, when




Figure 4. Three-person household – spatial trace.

joint activities are overemphasized by the utility function. At this point, high-utility individual activities are dropped from the schedule, which returns in a lower fitness level. It is difficult to generalize from this one experience, as the content of the calendar is likely to influence the results, especially if the schedule has exhausted the joint activities at some point. Nevertheless, this first experiment is reassuring, as it demonstrates that the algorithm responds in the expected way.

6. Discussion and further work

The work reported here outlines a method to find near-optimal schedules for households. It opens up new avenues for realizing the potential of activity-based approach in understanding and forecasting individuals as well as households. Obviously, many questions need to be addressed before this potential can be fully realized.



Figure 5. Three-person household – legend.

Computational performance needs to be improved before the household scheduler can be used in large-scale applications. On a Pentium IV 2.4 GHz Xeon system, a model run for the presented three-person example takes about 3 s per household member, which is acceptable for many purposes. Required calculation time rises with the size of the activity calendar and the number of generated trips. However, for iterative large-scale applications, such as multi-agent micro simulation (Raney and Nagel 2005) involving 10^6 persons, it is too slow. Strategies must be found to improve calculation speeds by at least a factor 100 to make integration possible without producing excessive run times for the combined model. While brute power (faster or more CPUs) helps, the solution has to be in smarter search strategies: parameterized rules to limit the number of generations, optimized parame-




Figure 6. Development of results.




Figure 7. Shared time and fitness as a function of in the three-person household scenario.

ters of the GA itself, using more than just one schedule from the generated populations, and/or recycling solutions as starting points for households facing only slightly different conditions (e.g. work place and home locations).

These steps must be taken to permit the estimation of the utility functions. Three approaches are possible for this calibration:

- Implementation of a second GA (or other heuristic optimization algorithm), which optimizes the parameters of the utility function for a given set of schedules. This avoids generation of the schedules, but raises the question whether these parameters could produce realistic new schedules.
- Generation of new schedules, comparison with a sample of observed schedules, and optimization of the utility function to maximize the similarity.
- Use of the household scheduler to generate good alternative schedules and employ a discrete choice approach to estimate the parameters. As the estimated parameters might be sensitive to the parameters chosen to generate the alternatives, iterations might be necessary.

Speed is essential, as comprehensive measurement of the fit between observed and simulated schedules requires substantial computing efforts in itself (Joh 2004). Even for models without random parameters, the schedules will have to be generated often before the parameter estimate converges. High speeds will be required to allow extensive tests of the utility functions' of different formulations.

At this point, the algorithm does not directly include monetary expenditures of activities at all. It is clear that this is desirable in the long term, especially if the model were to address scheduling over multiple days. Still, before this can be achieved, there is a need to improve our empirical understanding of the link between activity costs and hedonic benefit gained (see (Axhausen et al. 2002) for an example of relevant survey work).

Another important intra-household interaction can be seen in trips shared by household members. In the Mobidrive dataset, 18% of all trips were shared by at least two household members (Singhi 2001). The modelling of joint activities as presented here is a first necessary step to capture joint trips in activity-based analysis.

Acknowledgements

The paper is based on the M.Sc. thesis of the first author at the University Osnabrück, which he undertook at the Institute for Transport Planning and Systems, ETH Zurich (Meister 2004). The authors gratefully acknowledge the great contribution of David Charypar, who made the code of his single person scheduler available to us. It served as the starting point of our work.

References

- Arentze T & Timmermans H (2000) *ALBATROSS: A Learning-based Transportation Oriented Simulation*. Eindhoven: EIRASS.
- Axhausen KW (1990a) A simultaneous simulation of activity chains, . In Jones PM (ed.), *New Approaches in Dynamic and Activity-based Approaches to Travel Analysis* (pp. 206–225). Aldershot: Avebury.
- Axhausen KW (1990b) Judging the day: A synthesis of the literature on measuring the utility of activity patterns. Working Paper, 561, TSU, Oxford: University of Oxford.
- Axhausen KW (1996) The design of environmentally aware travel diaries. *Transportation Planning and Technology* 19(3): 275–290.
- Axhausen KW (October 2002) Some ideas for a microsimulation system of travel demand. Internal presentation, ETH Zürich.
- Axhausen KW & Goodwin PB (1991) EUROTOPP: Towards a dynamic and activity-based modelling framework. In: *Advanced Telematics in Road Transport* (pp.1020–1039). Amsterdam: Elsevier.
- Axhausen KW Zimmermann A Schönfelder S Rindsfuser G & Haupt T (2002) .Observing the rhythms of daily life: A six-week travel diary. *Transportation* 29(2): 95–124.
- Bates J Williams I Coombe D & Leather J (1996) .The London Congestion Charging Programmed: 4. The transport models. *Traffic Engineering and Control* 37(5): 334–339.
- Becker GS (1965) A theory of the allocation of time. *Economic Journal* 75 493–517.
- Bhat CR Frusti T Zhao H Schönfelder S & Axhausen KW (2004) .Intershipping duration: An analysis using multi-week data. *Transportation Research* 38B(1): 39–60.
- Bovy PHL & Stern E (1990) *Route Choice and Way Finding in Transport Networks*. Dordrecht: Kluwer.
- Bowman JL (1998) *The day activity schedule approach to travel demand analysis*, Ph.D. thesis. Cambridge, MA: Massachusetts Institute of Technology.
- Bowman JL, Bradley M, Shiftan Y, Lawton TK & Ben-Akiva ME (July 1998) Demonstration of activity-based model system for Portland. Paper presented at the 8th World Conference on Transport Research, Antwerp, Belgium.
- Charypar D & Nagel K (2003) Generating complete all-day activity plans with genetic algorithms. Paper presented at the 10th International Conference on Travel Behaviour Research, Lucerne, August 10–14. <http://www.sim.inf.ethz.ch/publications/ga-acts-iatbr03/>. Accessed April 22, 2005.
- Damm D (1983). Theory and empirical results a comparison of recent activity-based research, . In Carpenter S & Jones PM (Eds.), *Recent Advances in Travel Demand Analysis* (pp. 3–33). Aldershot: Gower.
- De Serpa AC (1971) A theory of the economics of time. *Economic Journal* 81 828–846.
- Doherty ST & Axhausen KW (1998). A unified framework for the development of a weekly household activity-travel scheduling model, . In Brilon W, Huber F, Schreckenberg M & Wallentowitz H (Eds.), *Traffic and Mobility* (pp. 35–56). Heidelberg: Springer.
- Doherty ST & Miller EJ (2000) .A computerized household activity scheduling survey. *Transportation* 27(1): 75–97.
- Ettema DF Borgers AWJ & Timmermans HJP (1997) .A simulation model of activity scheduling behaviour. *Transportation Research Record* 1413 1–11.
- Fellendorf M Haupt T Heidl U & Scherr W (1997). PTV Vision: Activity-based micro-simulation model for travel demand forecasting, . In Ettema DF & Timmermans HJP (Eds.), *Activity-based Approaches to Travel Analysis* (pp. 55–72). Oxford: Pergamon.
- Gärbling T, Brännäs K, Garvill J, Golledge RG, Gopal S, Holm E & Lindberg E (1989) Household activity scheduling. *Transport Policy, Management & Technology Towards 2001, IV* (pp. 235–248). Ventura: Western Periodicals.

- Gärling T Kalen T Romanus J & Selart M (1998) .Computer simulation of household activity scheduling. *Environment and Planning A* 30(4): 665–679.
- Gärling T Kwan M-P & Golledge RG (1994) .Computational-process modelling of household activity scheduling. *Transportation Research* 28B(5): 355–364.
- Gliebe JP & Koppelman FS (2002) .A model of joint activity participation. *Transportation* 29(1): 9–72.
- Goldberg DE (1989) *Genetic Algorithms in Search, Optimization and Machine Learning* Reading: Addison-Wesley.
- Jara-Diaz SR & Farah M (1988) .Valuation of users' benefits in transport systems. *Transport Reviews* 8(2): 197–218.
- Joh CH (2004) Measuring and predicting adaptation in multidimensional activity-travel patterns, bouwstenen, No. 79. Faculteit Bouwkunde, TU Eindhoven, Eindhoven.
- Jones PM (1979) HATS: A technique for investigating household decisions. *Environment & Planning A* 11(1): 59–70.
- Jones PM Dix MC Clarke MI & Heggie IG (1983) *Understanding Travel Behaviour*. Aldershot: Gower.
- Jones PM Koppelman F & Orfeuil JP (1990). Activity analysis: State-of-the-art and future directions, . In Jones PM (Eds.), *New Approaches in Dynamic and Activity-based Approaches to Travel Analysis* (pp. 34–55). Aldershot: Avebury.
- Kitamura R (June 1996) Applications of models of activity behavior for activity based demand forecasting. In: *TMIP (Travel Model Improvement Program) Activity-based Travel Forecasting Conference*. <http://tmip.fhwa.dot.gov/clearinghouse/docs/abtf>. Accessed April 22, 2005.
- Lynch K (1960) *The Image of the City* Cambridge: MIT Press.
- Meister K (2004) Erzeugung kompletter Aktivitätenpläne für Haushalte mit genetischen Algorithmen. Master thesis, Institute for Transportation Planning and Systems (IVT), ETH ZPürich, Zürich.
- Noland RB & Small KA (1995) .Travel time uncertainty, departure time choice and the cost of morning commutes. *Transportation Research Record* 1493 150–158.
- Pas EI & Harvey AS (1997). Time use research and travel demand analysis modelling, . In Stopher PR & Lee-Gosselin MEH (Eds.), *Understanding Travel Behaviour in an Era of Change* (pp. 315–338). Oxford: Pergamon.
- Poock M & Zumkeller D (1976) Die Anwendung einer massnahmenempfindlichen Prognosemethode am Beispiel des Grossraums Nürnberg. DVWG-Workshop Policy Sensitive Models, Giessen.
- Raney B & Nagel K (January 2005) An improved framework for large-scale multi-agent simulations of travel behavior. Paper presented at the 84th TRB Annual Meeting, Transportation Research Board, Washington DC.
- Recker W The household activity pattern problem: General formulation and solution, *Transportation Research B*, 1999; 29(1): 61–77.
- Rindsfuser G et al. (August 2003) Tracing the planning and executing attributes of activities and travel: Design and application of a hand-held scheduling probes survey. Paper presented at the 10th International Conference on Travel Behaviour Research, Lucerne.
- Roorda M & Miller EJ (May 2004) Strategies for resolving activity scheduling conflicts: An empirical analysis. Paper presented at the EIRASS conference, Maastricht.
- Schönfelder S & Axhausen KW (2001). Modelling the rhythms of travel using survival analysis, . In Kaspar C, Laesser C & Bieger T (Eds.), *Jahrbuch 2000/2001 Schweizerische Verkehrswirtschaft* (pp. 137–162). St. Gallen: SVWG.
- Schönfelder S & Axhausen KW (2003) .Activity spaces: Measures of social exclusion?. *Transportation Policy* 10(4): 273–286.
- Singhi P (2001) Analysis of joint trips using C++ in Mobidrive. *Arbeitsberichte Verkehrs- und Raumplanung*, No. 87, IVT, ETH Zürich, Zürich.

Zhang J, Fujiwara A, Timmermans H & Borgers A (May 2004) Methodology for modeling household time allocation behavior. Proceedings of the EIRASS Conference on Progress in Activity-based Analysis, Maastricht, The Netherlands.

About the authors

Prof. Dr. Kay W. Axhausen is the Professor of Transport Planning at the Swiss Federal Institute of Technology (ETH) in Zurich. He holds his post in the Institute for Transport Planning and Systems of the Department of Civil, Environmental and Geomatic Engineering. Before his appointment at ETH, he worked at the Leopold Franzens University of Innsbruck, Imperial College London and the University of Oxford. He holds a Ph.D. in Civil Engineering from the University of Karlsruhe and an M.Sc. from the University of Wisconsin -- Madison.

He has been involved in the measurement and modelling of travel behaviour for the last 20 years contributing especially to the literature on stated preferences, microsimulation of travel behaviour, valuation of travel time and its components, parking behaviour, activity scheduling and travel diary data collection. His current work focuses on the microsimulation of daily travel behaviour and long-term mobility choices and the response of the land-use system to those choices. This work is supported by analyses of human activity spaces and their dependence on the traveller's personal social network.

He is the chair of the International Association of Travel Behaviour Research (IATBR) and editor of *Transportation*.

Konrad Meister has been a member of the ETH Institute for Transportation Planning and Systems since September 2004. His work focuses on modelling travel demand for agent-based microsimulation of transport systems. He holds a diploma in Applied Systems Science from the University Osnabrück and a bachelor's engineering degree in Information Technology from the University of Cooperative Education Saxony in Dresden.

Martin Frick has been a member of the ETH Institute for Transportation Planning and Systems, since August 2002. His current work focuses on traffic microsimulation, monte carlo simulation, synthetic population generation and agent-based traffic demand modelling. He holds a diploma in physics from the University Konstanz.