STRC

A Framework to Represent Joint Decisions in a Multi-Agent Transport Simulation

Thibaut Dubernet

Kay W. Axhausen

STRC **April 2013**

STRC | 13th Swiss Transport Research Conference

Monte Verità / Ascona, April 24 – 26, 2013

STRC

A Framework to Represent Joint Decisions in a Multi-Agent Transport Simulation

Thibaut Dubernet
Institute for Transport Planning and Systems (IVT)
ETH Zurich
CH-8093 Zurich
phone: +41-44-633 68 65
fax: +41-44-633 10 57
thibaut.dubernet@ivt.baug.ethz.ch

Kay W. Axhausen
Institute for Transport Planning and Systems (IVT)
ETH Zurich
CH-8093 Zurich
phone: +41-44-633 39 43
fax: +41-44-633 10 57
axhausen@ivt.baug.ethz.ch

April 2013

Abstract

Mobility simulation softwares are tools which can be used for policy evaluation or to test behavioral assumptions.

While a wide range of softwares are being developped, with a wide range of approaches and assumptions, most of them focus on the representation of *individual* decision making. *Joint* decision making, on the contrary, has received much less consideration, probably due to the difficulty to represent it accurately and efficiently, and the reasonable predictions that can already be produced whithout considering it.

Joint decisions may however have an important influence on mobility behaviors: household members need to coordinate to use a limited number of vehicles; groups of friends need to negociate on a place to get dinner. In particular, it is hypothesised that location choice for leisure activities is strongly influenced by the possibility to meet social contacts at destination. Thus, representing accurately such processes may allow to fill the gap observed between predictions and observations in travel distance distribution for leisure trips.

This paper presents an approach to simulate such joint decision processes using the MATSim multi-agent simulation framework. It uses two main components: an approach to generate *joint plans*, and an approach to correlate individual-level plan choice, knowing the constraints imposed by the joint plans. The approach is designed to be usable on arbitrary social network

structures.

The behavior on a simple test scenario for the social network of intra-household ties is presented.

1 Introduction

Traffic simulation models are used to predict traffic flows on a network, aiming at supporting analysis and decision taking. Since the first models proposed in the middle of the twentieth century, the increase in computational power and the continuous improvement of traffic models led to always more precise and finer predictions. A popular framework for simulating traffic is multi-agent activity based transport simulation, where agents, representing individuals, travel through a simulated network from one activity to the other.

With such models, the simulation is based on individual behavioural models, and thus virtually allows to simulate any behaviour which impact on traffic is assumed, or known, to be important. The fact that several persons may coordinate with others is such a behavior. This kind of behavior for instance occurs when two or more individuals travel together, meet for a joint activity, or coordinate to make the best use of common resources, such as an household car.

MATSim is a multi-agent simulation software, which uses an evolutionnary algorithm to search for a user equilibrium. This purely competitive model is not, in its current state, adapted to the simulation of cooperative behavior. To solve this issue, we propose a way to include coordination behavior between arbitrary agents, without breaking the assumption that the behavior of (most of) the other agents is random from the agent's perspective.

This paper presents this approach in detail. First, related work is reviewed in Section 2. The approach, as well as its implementation as a MATSim module, are presented in Section 3. The results of a simulation of intra-household ride sharing for a simple test scenario are presented in Section 4.

2 Related studies

2.1 The activity-based approach

Simulation of travel behaviour is a widely used tool, which can be used for predicting the effects of some change in infrastructure, reconstruct missing data about the current state, policy evaluation or behavioural hypothesis testing.

Generally, simulation models are classified as macroscopic, mesoscopic or microscopic, depending on the level of aggregation used.

Naturally, each type of model has its strengths and weaknesses. While macroscopic models, which only work at the aggregated level, are computationally efficient and only require aggregated data as input, they have difficulties to represent time-varying aspects of traffic. On the other hand, microscopic models, by simulating agents individually, can predict traffic dynamics much more easily, but at a much higher computational cost and with finer data as input.

However, increase in computational power in the last decades has made this kind of models more and more appealing.

An appealing framework while simulating individuals at a disaggregated level is to use the so-called "activity-based " approach, proposed during the early eighties (Jones et al. (1983), Recker et al. (1986)). In this approach, the fact that travel is always oriented toward a goal is taken into account explicitly: agents are assigned plans, consisting of located activities, and travel between those activities in a simulated network. A fundamental difference with trip-based approaches is the explicit modeling of travel as a need *derived* from the need or willingness to perform activities (McNally and Rindt (2008)).

The way the plans are computed depends on the model: in the following, we focus on the way the MATSim software achieves this task.

2.1.1 Equilibrium based models: the MATSim process

MATSim is an open-source software, released under the terms of the GNU Global Public License (GPL). It aims at simulating time-dependant mobility flows (Balmer et al. (2008), Rieser et al. (2007)). To do so, it relies on the assumption that the state of traffic on an average day corresponds to a user equilibrium: no individual can improve the utility he gets from his day by modifying his daily plan, given the plans of the rest of the population. The only dimensions considered in the equilibrium are the ones related to short term decisions: route choice, mode choice, departure times, *etc.*.

More formally, finding the equilibrium consists in solving

$$\max_{p_i \in P_i} U(p_i \mid p_{-i}) \tag{1}$$

for each agent i, where p_i is the plan of agent i, P_i the set of possible plans for agent i, and p_{-i} the set of the plans of other agents.

To search for such an equilibrium, MATSim uses a co-evolutionary process, where each agent performs an evolutionary algorithm to solve the problem (1). The process is as follows: starting with initial plans, agents are moved through a simulated network, giving estimates of the cost of travel (and thus of the influence of p_{-i}). Then, plans of a given fraction of the agents are copied and mutated, randomly or to optimality given the previous state. Non mutated agents choose one of their previous plans based on the past scores, and the simulation is run again. If the number of plans of a agent is over a fixed memory size, the worst plan of this agent is discarded, pushing the evolution toward better plans. This process is iterated until a stopping criterion is met (currently, a fixed number of iterations fixed *a priori* is used (Meister et al. (2010))).

This process allows to take into account the complex relationship between congestion and individual's daily plans. It can be considered both as an algorithm to search for a user equilibrium or as an actual simulation of human learning (Nagel and Marchal (2006)).

Currently, replanning can include least-cost re-routing, location choice (Horni et al. (2009)), duration and mode optimisation (Meister et al. (2006)). Experiments have also included activity sequence (Feil (2010)).

2.1.2 Using optimisation algorithms to replan agents

As pointed out before, the relaxation process consists in iteratively improving the plans of the agents, knowing the previous behaviour of other agents, until a steady state is reached.

The standard approach, based on evolutionnary algorithm and Machine Learning, consists in randomly "mutating" some plans between iterations. Each agents possesses a memory, which stores a fixed number of past plans, allowing to revert changes implying a decrease in utility.

However, another approach has been implemented since then, making use of optimisation algorithms in the mutation step. This approach was shown to allow MATSim to converge in fewer iterations to an equilibrium state, with a score at least as high as with random mutation. Optimisation algorithms used include least-cost routing, activity duration optimisation with CMA-ES (Charypar et al. (2006)) or genetic algorithm (Meister et al. (2006)), or activity sequence and other properties with Tabu Search (Feil et al. (2009)).

This approach however tends to result in less diversity in the agent's plans, and should be used with care.

2.2 Joint decisions modeling

The random utility theory is a well-known and extensively studied way of predicting individual's behaviour, which is widely used in transportation research (Ben-Akiva and Lerman (1985)). In this general framework, each alternative is associated a numerical utility, composed of a systematic part (its expectation) and a random error term (representing unobserved variability). The probability for an individual to choose one of the alternatives corresponds to the probability for the utility of this alternative to be higher than the utility of all other alternatives.

This framework as been applied to joint decision making, and to joint scheduling in particular: we provide here a review of those studies.

Aside from these random utility models , non-probabilistic utility maximisation techniques have been proposed for creating schedules for activity based transport simulation: we present here some of those attempts for household plans generation.

2.2.1 Random utility based models

The random utility theory as been applied early to joint decision modeling, by considering the choice problem as a group utility maximisation problem.

In the last decades, this framework began to be applied to group (mainly household) schedules generation for activity based transport simulation.

However, the choice set is of high dimension, with both discrete (activity types, joint activity participation, sequence of activities, modes *etc.*) and continuous (activity duration) dimensions. Thus, depending on the authors, different choice dimensions are considered.

Zhang, Timmermans and Borgers develop a model where time for different activity types is allocated to household members, subject to time constraints (including equality of time participation in joint activities) (Zhang et al. (2005)). Given individual random utilities for the different activity type, their model gives deterministic time allocation.

Bradley and Vovsha focus on the "daily activity pattern" generation, with household "maintenance" tasks (*e.g.* shopping) allocation and possibility of joint activities (Bradley and Vovsha (2005)). To do so, they assume a layered choice structure: first, a daily activity pattern is assigned to household members; then, "episodic" joint activities can be generated; finally, maintenance activities are assigned.

Gliebe and Koppelmann (Gliebe and Koppelman (2005)) also base their models on the daily activity pattern concept. In their model, the joint outcome (the succession of individual and joint activities) is first determined, and individuals then choose an individual pattern compatible with the joint outcome. The same authors also derived a constrained time allocation model, which predicts the time passed by two individuals in joint activities (Gliebe and Koppelman (2002)). Rather than postulating a group-level utility function, those models specify a special distribution for the error terms of the individuals. In this setting, the error term of the individuals are correlated so that the probability of choosing a given joint output is the same for all individuals.

Miller, Roorda and Carrasco develop a model of household mode choice (Miller et al. (2005)). The main difference with an individual mode choice model is the consideration of household-level vehicle allocation. In their model, individuals first choose modes individually. If a conflict occur, the allocation that maximizes the household level utility is chosen. The members which were not allocated the vehicle will report on their second best choice, and/or examine shared rides options.

2.2.2 Alternative approaches

Aside from the random utility theory based models, some other ways to deal with joint scheduling have been proposed.

Golob and McNally propose a structural equation model, which predicts time allocation and trip chaining based on descriptive variables of an household (Golob and McNally (1997)). Golob also used this structural equation model approach to model the dependency of time allocations of the two heads (man and woman) of an household (Golob (2000)).

Another class of approaches is the use of optimisation algorithms to generate households plans. They handle the household scheduling problem by transforming it into a deterministic utility maximisation problem. Contrary to the previously presented approaches, those alternatives did not lead to estimate a model against data.

The first of those approaches was proposed by Recker in the mid nineties (Recker (1995)). By extending increasingly the formulation of the Pick-Up and Delivery Problem With Time Windows, which is a well studied combinatorial optimisation problem, he formulates the problem of optimising the activity sequence of members of an household as a mathematical programming problem, taking into account vehicle constraints, individual and household level activity, possibility of choosing whether to perform or not an activity, with the possibility of shared rides. However, due to the complexity of the problem, the full problem cannot be solved exactly by standard operations research algorithms, and the activity durations are not part of the optimised dimensions. However, Chow and Recker (2012) designed an inverse optimisation method to calibrate the parameters of this model, including the time window constraints, using measured data. Also, the formulation from Recker (1995) was latter extended by Gan and Recker (2008) to introduce the effects of within-day rescheduling due to unexpected events.

A more recent attempt to generate plans for households uses a genetic algorithm, building on a previous genetic algorithm for individual plan generation (Charypar and Nagel (2005), Meister et al. (2005)). This algorithm optimises sequence, duration and activity choice for an household, rewarding the fact for several members of the household to perform the same activity simultaneously (*i.e.* "joint activities").

Finally, rule-based systems have been developed, which use heuristic rules to construct household plans. They often rely on a simulated bargaining process (Arentze and Timmermans (2009), Ma et al. (2011, 2012)).

3 Inclusion of coordination in a multi-agent transport simulation

As stated previously in Section 2.1.1, the equilibrium approach to activity-based travel modeling consists in assuming that individuals, represented as software agents, attempt to maximize their utility given the state of the transport system, itself influenced by the behavior of other agents.

Our claim here is that this model is sometimes not sufficient: there are important cases when an individual must not only be aware of the global properties of the system as resulting from the behavior of others, but also of relevant parts of other's plans. That is, it is sometimes useful to consider coordination behaviors.

This section presents an adaptation of the evolutionary approach used in MATSim, well suited to simulate such behaviors.

3.1 Representing coordination in MATSim

Recently, various ways to cope with individual coordination have been proposed in the literature. Some of them rely on the actual simulation of bargaining processes, while other consider a utility-based optimisation of a *joint plan* (*cf.* Section 2.2).

The former is easier to generalize to arbitrary social structures , while the latter, by being formulated as a mathematical problem, is easier to interpret and better suited to calibration.

We propose here a utility-based approach based on the *joint plan* concept, but allowing to represent coordination in arbitrary social structures. As the usual MATSim process, it uses an evolutionnary algorithm to optimize full daily plans, that the agents follow "blindly" in a mobility simulation to obtain scores. To take into account coordination between agents, it relies on two complementary constraints on the possibility to select an individual plan, given the plans selected by relevant other agents:

1. **joint plans:** a joint plan is a set of individual plans of several agents, which must always be selected together. An example is the plan of a driver and his passenger, or plans of two household members using the same vehicle.
2. **incompatible (joint) plans:** There are also cases when one does not want two or more (joint) plans to be selected at the same time: for instance, two joint plans representing coordination for the usage of the same vehicle should not be selected at the same time.

3.2 Plan adaptation with coordination behaviors

Given those constraints on which plans of each agent's memory can be chosen, one needs a way to create new (joint) plans from old ones, and a way to select past plans based on the experienced score.

To achieve this, it is not possible anymore to consider agents in isolation, and one has to identify groups of agents to replan jointly. Fig. 1 illustrates the process to identify agents to replan together. In this figure, circles represent agents. Solid lines represent the existence of joint plans between agents, while discontinuous lines represent "social ties", that is, the possibility to create new joint plans. For replanning, agents having joint plans are put in the same group; while agents being linked by social ties can, but must not necessarily, be put in the same group. The groups used in different iterations need not be the same, as long as the constraints are respected.

Once groups are identified, the process is similar to the individual case: the first, mandatory step, is to select an individual plan for each agent in the group, considering the constraints. This is done by selecting the feasible combination of individual plans which maximizes the sum of some weights allocated to plans. Depending on the case, those weights can be the scores, randomized scores (to obtain the analogous of a random choice model), or be completely random (in order

Figure 1: Group identification

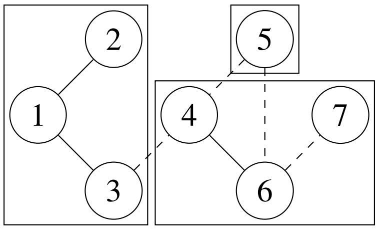

to draw a feasible combination at random). Due to the complexity of the problem, complete enumeration is impossible. Thus, a branch-and-bound approach is used to search for the best combination.

As in the individual case, the number of plans an agent can remember is limited. To push the evolution towards best plans, the plans pertaining to the feasible combination of worst score sum is removed, by taking care not to create states where no feasible combination remains.

The second, optional step, is to create new individual plans, by modifying a copy of the selected plan. All the strategies used at the individual level can of course be applied on the individual plans; group level strategies, which modify the joint plan structure (for instance by creating new shared rides) or act on joint plans (for instance by synchronizing plans of co-travelers) can also be applied.

4 Performance

4.1 Scenario

This section demonstrates the behavior of the framework on a simple scenario where agents in cliques (*e.g.* households) are allowed to perform two-persons ride sharing.

The network consists in a "home" and a "work" location, with a single road in between, with enough capacity to avoid congestion. Work is always open, so that agents can go any time of the day, to be sure that the coordination necessary to travel together comes from the joint decision framework rather than from the opening hours.

1022 agents are grouped in even-sized cliques of 2 to 50 members, half of the agents having no driver's license (*i.e.* the car mode is forbiden). There is the same number of clique of each size (plus or minus one). Agents start with home-work-home chains, with all trips made by walk. Public transport takes twice the time needed by car. All modes have the same disutility of travel time. In particular, this means that driving alone, driving somebody or being driven by somebody have the same disutility.

The replanning step is run by always taking the full cliques as replanning groups. The strategies executed and their probability are listed in Table 1.

Table 1: Probability of replanning strategies

Strategy	Description	Probability
Departure Time Mutation	Random mutation of activity end times	0.1
Mutation of Joint Trips	Grouping of random trips into joint trips, or random deletion of existing joint trips. After this process, activity end times are optimized, given the travel times observed in the previous iteration, in the best-response fashion described in Section 2.1.2.	0.2
Mode Mutation	Random allocation of a new mode to subtours	0.1
Logit-like Plan Selection	Selection of the feasible plan combination of maximum sum of randomized scores, using Gumbel-distributed error terms	0.6

Plans corresponding to a same joint trip are grouped in the same joint plans. There are no incompatible (joint) plans: any combination which respects the joint plans constraints is considered as valid.

The advantage of using such a simple scenario is that one can anticipate the results that should be expected, in order to check that the approach behaves as expected before using it on real-world complex scenarios. To assess the validity of the approach, the results of the run described above are compared with the results of a run with the same characteristics, but where all agents have access to a car. The expected results are that:

- that the share of individual vehicle mode (car, driver, passenger) be roughly equivalent with the case where all agents have access to a car. A different result would indicate a problem in the approach, such as a failure to coordinate agent's plans resulting in long waiting times.
- that plans containing passenger mode use this mode for both their trips, as it makes car unavailable to them for the other trip if they have a car, or is better than any other possible chain if they have no car.

As stated in Table 1, when inserting or deleting joint trips, a best response module, of the kind described in Section 2.1.2, is used to optimize activity durations and synchronize co-traveler plans. This has an influence on the diversity of the plans in the agent's choice set, and may have an influence on the results. Such modules are moreover more difficult to maintain and extend than combination of simple knowledge-free mutation modules. Thus, the results of this approach are compared to the results obtained when using a simpler approach, in which departure times of passengers are adapted to the ones of the driver.

4.2 Results

Table 2 shows the mode shares in the final iteration for the three scenarios. The share of individual motorized modes is roughly equivalent in both scenarios with best response, being slightly higher when all agents are allowed to drive a car. This indicates that the process is able to converge toward the expected state. Not using optimisation to synchronize plans results in a significantly lower share of individual motorized transport: by optimizing plans in which joint trips are inserted or removed, the individual motorized modes have more chance to be in a plan with an optimal time allocation than the other modes. A simple mutation approach results in more diversity in the plans, and is probably the approach to be favoured.

Table 2: Final mode shares

	Mode Shares (%)		
Mode	Half car (opt)	Half car (no opt)	All have car (opt)
Car	3.69	2.42	54.20
Car Driver	44.62	38.71	19.99
Car Passenger	44.62	38.71	19.99
Total Indiv. Veh.	*92.92*	*79.85*	*94.17*
Public Transport	6.94	18.91	2.08
Walk	0.07	0.40	1.91
Bike	0.07	0.83	1.85

Fig. 2 shows the evolution of the average minimum, executed and maximum joint plans size of the agents memory with iterations, in the two scenarios where half the agents only have a car. Note that joint plans of any size between 1 and the clique size are theoretically feasible: if an agent A drives and agent B in the morning and another agent C in the afternoon, the three plans have to be grouped in a same joint plan. The process naturally however converges toward small joint plans, after having generated bigger plans. The process without optimization takes more iterations to converge, and converges towards bigger joint plans on average. The reason is probably that when optimizing duration, agents having the same plan and utility function, when optimizing, the activities of co-travelers immediately start and end at the same time, making the creation of a second joint trip easy. When mutating, the heuristic synchronization rule may worsen the utility of performing an activity by getting it further from the optimal duration. This shows a difficulty of the current mutation approach to find the "obvious" solutions for the party composition.

Figure 2: Evolution of the average joint plan size per agent

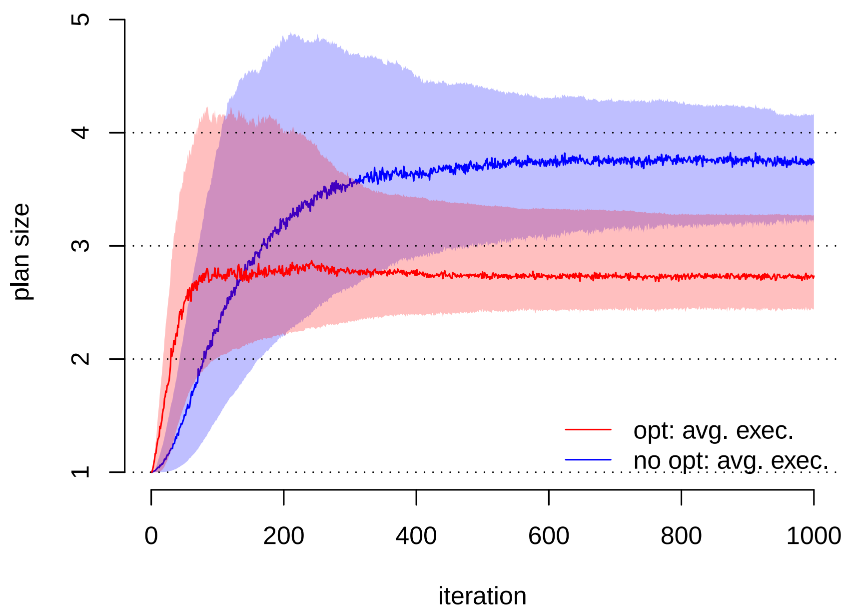

Table 3 shows the share of individual plans containing a given number of trips of a joint mode . As expected, for the simulation in which all agents have a car available, if an agent is a passenger, he is a passenger for his two trips. When half the agents have no car available, a significant number of agents have only one passenger trip: this probably comes from the lower number of possible better chains for agents without car (only passenger-passenger is better) compared to agents with car. Due to the fact that plans are selected using independent random error terms, the probability to select a chain increases with the number of occurrences of the chain in the choice set (the well-known "red-bus/blue-bus paradox" (Ben-Akiva and Lerman (1985))). This is a potential problem, common to all MATSim applications.

Table 3: Tour structure

		Share (%)		
Scenario	Mode	0 trip	1 trip	2 trips
Half cars	Either	1.6	18.4	80.0
	Driver	51.8	7.2	41.0
	Passenger	49.8	11.2	39.0
All cars	Either	48.3	23.8	27.9
	Driver	68.3	23.5	8.2
	Passenger	79.9	0.3	19.8

5 Conclusion

This paper presents a framework for simulating individual coordination in a multi-agent activity-based travel simulation software, and analyses the behavior of an implementation for the MATSim software.

The analysis shows that the framework is able to converge toward the expected state on a simple scenario for ride-sharing. The effect of using or not a best-response approach is also investigated. While a pure random mutation approach leads to more diversity of the plans in the agents' "choice sets", which is good, it seems to have difficulties to find the "obvious" solutions in term of party composition, namely two-persons car pools.

Future work includes investigating further the convergence issue of the mutation approach, and extend the approach to the case of coordination for the use of limited vehicle resources.

Acknowledgments

The work presented in this paper is part of the project "EUNOIA: Evolutive User-centric Networks for Intra-urban Accessibility", funded under the European Union's Seventh Framework Programme ICT Programme.

6 References

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.

Balmer, M., K. Meister, M. Rieser, K. Nagel and K. W. Axhausen (2008) Agent-based simulation of travel demand: Structure and computational performance of MATSim-T, paper presented at the *Innovations in Travel Modeling (ITM'08)*, Portland, June 2008.

Ben-Akiva, M. E. and S. R. Lerman (1985) *Discrete Choice Analysis: Theory and Application to Travel Demand*, MIT Press, Cambridge.

Bradley, M. A. and P. Vovsha (2005) A model for joint choice of daily activity pattern types of household members, *Transportation*, **32** (5) 545–571.

Charypar, D., K. W. Axhausen and K. Nagel (2006) Implementing activity-based models: Accelerating the replanning process of agents using an evolution strategy, paper presented at the *11th International Conference on Travel Behaviour Research (IATBR)*, Kyoto, August 2006.

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

Chow, J. Y. and W. W. Recker (2012) Inverse optimization with endogenous arrival time constraints to calibrate the household activity pattern problem, *Transportation Research Part B: Methodological*, **46** (3) 463 – 479.

Feil, M. (2010) Choosing the daily schedule: Expanding activity-based travel demand modelling, Ph.D. Thesis, ETH Zurich, Zurich.

Feil, M., M. Balmer and K. W. Axhausen (2009) New approaches to generating comprehensive all-day activity-travel schedules, *Research Report*, **575**, IVT, ETH Zurich, Zurich.

Gan, L. P. and W. W. Recker (2008) A mathematical programming formulation of the household activity rescheduling problem, *Transportation Research Part B: Methodological*, **42** (6) 571–606.

Gliebe, J. P. and F. S. Koppelman (2002) A model of joint activity participation, *Transportation*, **29** (1) 49–72.

Gliebe, J. P. and F. S. Koppelman (2005) Modeling household activity-travel interactions as parallel constrained choices, *Transportation*, **32** (5) 449–471.

Golob, T. F. (2000) A simultaneous model of household activity participation and trip chain generation, *Transportation Research Part B: Methodological*, **34** (5) 355–376.

Golob, T. F. and M. G. McNally (1997) A model of activity participation and travel interactions between household heads, *Transportation Research Part B: Methodological*, **31** (3) 177–194.

Horni, A., D. M. Scott, M. Balmer and K. W. Axhausen (2009) Location choice modeling for shopping and leisure activities with MATSim: Combining micro-simulation and time geography, paper presented at the *88th Annual Meeting of the Transportation Research Board*, Washington, D.C., January 2009.

Jones, P. M., M. C. Dix, M. I. Clarke and I. G. Heggie (1983) *Understanding Travel Behaviour*, Gower, Aldershot.

Ma, H., T. A. Arentze and H. J. P. Timmermans (2012) Incorporating selfishness and altruism int dynamic joint activity-travel scheduling, paper presented at the *13th International Conference on Travel Behaviour Research (IATBR)*, Toronto, July 2012.

Ma, H., N. Ronald, T. A. Arentze and H. J. P. Timmermans (2011) New credit mechanism for semicooperative agent-mediated joint activity-travel scheduling, *Transportation Research Record*, **2230**, 104–110.

McNally, M. G. and C. R. Rindt (2008) The activity-based approach, *Working Paper*, **UCI-ITS-AS-WP-07-1**, Institute of Transportation Studies, University of California, Irvine.

Meister, K., M. Balmer, K. W. Axhausen and K. Nagel (2006) planomat: A comprehensive scheduler for a large-scale multi-agent transportation simulation, paper presented at the *11th International Conference on Travel Behaviour Research (IATBR)*, Kyoto, August 2006.

Meister, K., M. Balmer, F. Ciari, A. Horni, M. Rieser, R. A. Waraich and K. W. Axhausen (2010) Large-scale agent-based travel demand optimization applied to Switzerland, including mode choice, paper presented at the *12th World Conference on Transportation Research*, Lisbon, July 2010.

Meister, K., M. Frick and K. W. Axhausen (2005) A GA-based household scheduler, *Transportation*, **32** (5) 473–494.

Miller, E. J., M. J. Roorda and J. A. Carrasco (2005) A tour-based model of travel mode choice, *Transportation*, **32** (4) 399–422.

Nagel, K. and F. Marchal (2006) Computational methods for multi-agent simulations of travel behavior, in K. W. Axhausen (ed.) *Moving Through Nets: The Physical and Social Dimensions of Travel*, 131–188, Elsevier, Oxford.

Recker, W. W. (1995) The household activity pattern problem: General formulation and solution, *Transportation Research Part B: Methodological*, **29** (1) 61–77.

Recker, W. W., M. G. McNally and G. S. Root (1986) A model of complex travel behavior: Part I - theoretical development, *Transportation Research Part A: Policy and Practice*, **20** (4) 307–318.

Rieser, M., K. Nagel, U. Beuck, M. Balmer and J. Rümenapp (2007) Truly agent-oriented coupling of an activity-based demand generation with a multi-agent traffic simulation, paper presented at the *86th Annual Meeting of the Transportation Research Board*, Washington, D.C., January 2007.

Zhang, J., H. J. P. Timmermans and A. W. J. Borgers (2005) A model of household task allocation and time use, *Transportation Research Part B: Methodological*, **39** (1) 81–95.