Including joint trips in a multi-agent transport simulation

Conference Paper

Author(s):
Axhausen, Kay W. ; Dubernet, Thibaut Jean Pierre

Publication date:
2012

Permanent link:
https://doi.org/10.3929/ethz-a-007250989

Originally published in:
Arbeitsberichte Verkehrs- und Raumplanung 764

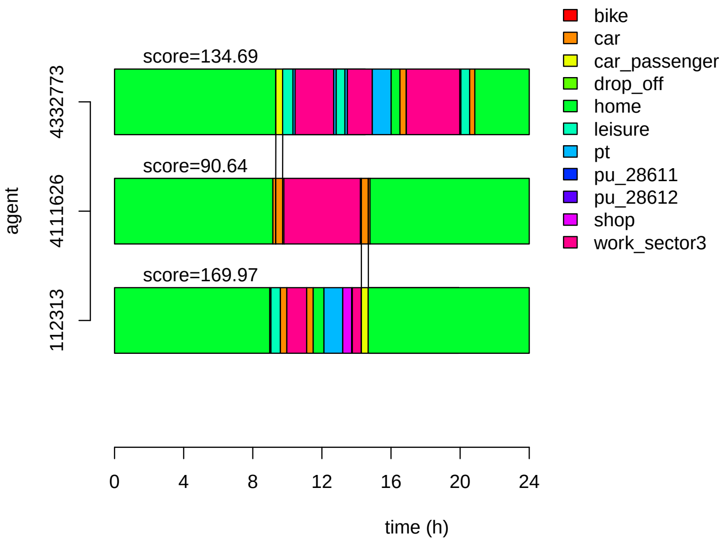

Including joint trips in a multi-agent transport simulation

Thibaut Dubernet

Kay W. Axhausen

STRC 2012 **May 2012**

STRC

12th Swiss Transport Research Conference
Monte Verità / Ascona, May 2 – 4, 2012

STRC 2012

Including joint trips 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:
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

May 2012

Abstract

Among the simulation approaches, the activity-based approach is used increasingly in operational models. It allows both to focus on individual behavioural rules, and to model behaviour at a much finer level than more aggregated approaches.

A behaviour that one can and should model is the *joint travel* behaviour: the fact that several individuals may travel in the same private vehicle. An appropriate modeling of such a behaviour is important both for detailed simulations of households behaviour and for the evaluation of some policies, like incentives to perform car-pooling.

We present an approach to include such a simulation of joint travel in a multi-agent microsimulation, considering the case of MATSim as an example.

First results of this approach for the Zürich region are presented.

Keywords
activity-based, joint trips, car pooling, microsimulation

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 successful 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 themselves to travel together is such a behaviour. Its simulation can be used to predict the impact of car-pooling on traffic flows, or the impact of car-pooling incentives on the share of this mode. Another major application is to better simulate the behaviour at the household level, where the share of such trips is known to be significant.

MATSim is a multi-agent simulation software, which uses an Evolutionnary Algorithm to search for a user equilibrium. This purely competitive model is not adapted to the simulation of cooperative behavior. To solve this issue, we propose to define the equilibrium over competitive cliques, defined as groups of agents aiming at maximising a group-level utility.

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 run of a scenario for the urban area of Zürich, Switzerland, 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.

A successful 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 mainly aims at simulating time-dependant traffic 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.*.

To search for such an equilibrium, MATSim uses a relaxation process: starting with initial plans, agents are moved through a simulated network, giving estimates of the cost of travel. Then, plans of a given fraction of the agents are 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. This process is executed 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 traffic flows and the utility of a plan. 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)). Depending on the approach, the way replanning will be handled will be slightly different: if one searches for a Nash equilibrium, each agent should have an optimal strategy at the end of the process; if one searches to simulate human learning, sub-optimal strategies may be allowed (and should appear), as soon as they result from a behaviourally sound search process.

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 Plan scores in MATSim

As pointed out, MATSim mainly consists in a daily plan optimisation algorithm coupled to a traffic flow simulation. Thus, a performance metric is needed.

In general, MATSim uses the so-called "Charypar-Nagel" scoring function, first introduced to generate daily plans out of the iterative MATSim process (Charypar and Nagel (2005)).

In this formulation, the utility of a plan takes the form of a sum of the activity of performing activities and of the disutility of traveling:

$$F = \sum_{i=1}^{n} U_{act}(\mathrm{type}_i, \mathrm{start}_i, \mathrm{dur}_i) + \sum_{i=2}^{n} U_{trav}(\mathrm{loc}_{i-1}, \mathrm{loc}_i) \tag{1}$$

where the utility of an activity is:

$$U_{act,i} = U_{dur,i} + U_{wait,i} + U_{late.ar,i} + U_{early.dp,i} + U_{short.dur,i} \tag{2}$$

and:

$$U_{dur}(t_{dur}) = \beta_{dur}\, t^{*} \ln\left(\frac{t_{dur}}{t_0}\right) \tag{3a}$$

$$U_{trav}(t_{trav}) = \beta_{trav}\, t_{trav} \tag{3b}$$

$$U_{wait}(t_{wait}) = \beta_{wait}\, t_{wait} \tag{3c}$$

$$U_{late.ar}(t_{start}) = \begin{cases} \beta_{late.ar}(t_{start} - t_{latest.ar}) & \text{if} t_{start} > t_{latest.ar} \\ 0 & \text{otherwise} \end{cases} \tag{3d}$$

$$U_{early.dep}(t_{end}) = \begin{cases} \beta_{early.dep}(t_{earliest.dep} - t_{end}) & \text{if} t_{end} < t_{earliest.dep} \\ 0 & \text{otherwise} \end{cases} \tag{3e}$$

$$U_{short.dur}(t_{end}) = \begin{cases} \beta_{short.dur}(t_{short.dur} - (t_{end} - t_{start}) & \text{if} t_{end} < t_{short.dur} \\ 0 & \text{otherwise} \end{cases} \tag{3f}$$

where:

t^{*}	is the typical duration for the activity
t_0	is the minimal duration for the activity
t_{dur}	is the actual utility duration
t_{trav}	is the traveling time
t_{wait}	is the waiting time
t_{start}	is the start time
t_{end}	is the end time

2.1.3 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)).

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").

3 Inclusion of joint trips in a multi-agent simulation

3.1 Joint trips and user equilibrium

In an equilibrium-based traffic microsimulation, one searches the equilibrium state where no agent can improve its utility by modifying its individual strategy, given the strategies of the other agents. By strategy, one usually understands “daily plan”, that is, sequence, location and duration of activities, modes and routes. The utility of the strategy of an agent is influenced by the strategy of other agents via congestion only. Thus, the effect of an individual change of strategy (*e.g.* a change in the departure time of a single agent) is likely to have only minor influence on the utility of other agents’ plans.

This situation is somewhat changed when including the possibility for several agents to travel in the same private vehicle: the utilities that the co-travelers get from their daily plans is highly dependant on the plans of the other co-travelers. For example, the score of a passenger’s plan is highly dependant on whether the driver actually chooses to pick him up or not.

Another phenomenon changes significantly the situation: the utility of an agent's plan may depend on the utility of the plans of his co-travelers, that is, "utility transfers" may occur. Those utility transfers may have different causes: willingness to help (*e.g.* "serve passenger" rides in households), monetary compensations (in the case of car-pooling services) *etc.*.

As seen in Section 2.2, a classic way of dealing with this kind of issues is to consider a group-level utility, made by aggregating individual utilities. Members of the group are assumed to be willing to maximise the utility of the group rather than their own.

Even though we follow this tradition, the specification (and the existence) of this group utility is problematic. Particularly, it seems likely that no such function can be defined for totally unrelated people: identifying collaborating agents is a challenging task.

Finally, if the choice of co-travelers is included in the search process, the range of variables taken into account into the utility function should be enlarged. For instance, social closeness is likely to have a predominant influence on both the possible joint trips available to a particular agent, and to the utility this agent will get from his trip.

For the current work, we do not focus on the problem of identifying which agents should travel together, but on the following: *given* that two or more agents may travel together, how to optimise their plans?

Considering this, when simulating joint trips, the equilibrium is defined over groups of agents. For this, we include the concept of *clique* in the multi-agent framework:

Definition 1 (Clique) *The set of agents is partitioned in **cliques**. Each agent pertains to one and only one clique. **Joint plans** are defined at the clique level.*

A clique may represent an household, a group of colleges, or whatever group of interest.

Definition 2 (Joint plan) *A **joint plan** is a set of individual plans, one for each agent of a clique. All individual plans of a joint plan are always chosen together. A unique score is affected to a joint plan.*

Please note that due to the transfer assumption, some care must be taken in the clique definition. In particular, cliques should be small enough so that travel times on the links cannot be changed dramatically by replanning just one clique. In the limit case with one "clique" containing all agents, the "equilibrium" would be the so-called "social optimum", which is different of the user equilibrium, as is well known (Roughgarden and Tardos (2002)).

3.2 MATSim implementation

As stated in Section 2.1.1, MATSim uses an iterative relaxation process to search for a user equilibrium. In this framework, agents' plans are iteratively simulated, scored and modified until no improvement can be made.

In MATSim, a daily plan is represented as a sequence of activities and "legs" (movement with one transport mode). In this context, a "joint trip" is a set of individual trips, themselves consisting in a sequence of legs, pick-up and drop-off activities.

To simulate joint trips, special care is needed at the replanning step, so that synchronized plans are generated. Thus, a special module, based on Tabu Search, takes care of optimizing activity durations and modes.

Time and mode optimisation At each iteration, for a small fraction of the cliques, activity durations and mode at the subtour level are optimized, based on time dependent travel time estimates based on the observed travel times in the previous run of the traffic flow simulation. This "best response" approach builds on previous successful experiments with the optimisation of various dimensions of an individual plan (Meister *et al.* (2006), Feil *et al.* (2009)).

Activity durations and modes are optimized using a Tabu Search algorithm, inspired by a previous approach for individual plans (Feil (2010)). Tabu Search is a metaheuristic method, initially aimed at solving combinatorial optimisation problems (Glover (1989)). It directs local search heuristics by the use of a *tabu list*, containing information on the previous moves. The rationale of this method is to avoid getting trapped in local optima, by preventing to re-explore already known parts of the search space. The interested reader may consult Glover and Taillard (1993) for a good introduction.

The details the algorithm are the following:

- **activity durations**:
  - **neighborhood**: each solution of the neighborhood corresponds to the current solution, where the duration of one single activity is increased or decreased by a given amount. A set of amounts is defined as a parameter of the algorithm.
  - **tabu list**: when such a move is selected, all moves going in the opposite direction (*e.g.* the moves decreasing the duration of an activity which duration was increased) are tabu for a fixed number of iterations.
- **modes**:

  - **neighborhood**: each solution of the neighborhood corresponds to the current solution, where the mode of a single subtour is changed.
  - **tabu list**: when such a move is selected, the deselected mode is tabu for a fixed number of iterations.

The different pick-ups related to a same joint trip must be synchronised. This synchronisation is enforced by penalising unsynchronised plans, with a penalty increasing linearly with the difference in end times.

4 First results

As an example of the output of the process, we present here the results for runs on the urban area of Zürich, Switzerland. The runs used a population of agents containing 10% of the real population, and used the default values for the parameters of the utility. Only the *car* mode suffers from congestion. In particular, passengers are moved according to the expected travel time, rather than with the driver.

4.1 Cliques and joint trips

In its current state, the implementation of the model works only with pre-identified possible joint trips. In the current case, those possible joint trips were generated by a partner, using a software aiming at generating possible car-pooling matches. Based on the output of a run without joint traveling, it attached passengers to drivers using a maximum detour with time windows approach.

All passengers traveling together were grouped in a same clique. It may be worth noting that clique sizes can get quite large (for example, two agent A and B being both driven separately by an agent C must pertain to the same clique, even if they never travel together). For computing time reasons, only cliques of less than 10 members were kept in the sample. Figure 1 shows an histogram of cliques sizes.

4.2 Running time

Table 1 presents the probability of the replanning modules, and Table 2 presents statistics relative to running time. The "replanning" time includes all replanning modules. The run was performed

Figure 1: Number of cliques per clique size (cliques of more than 2 members)

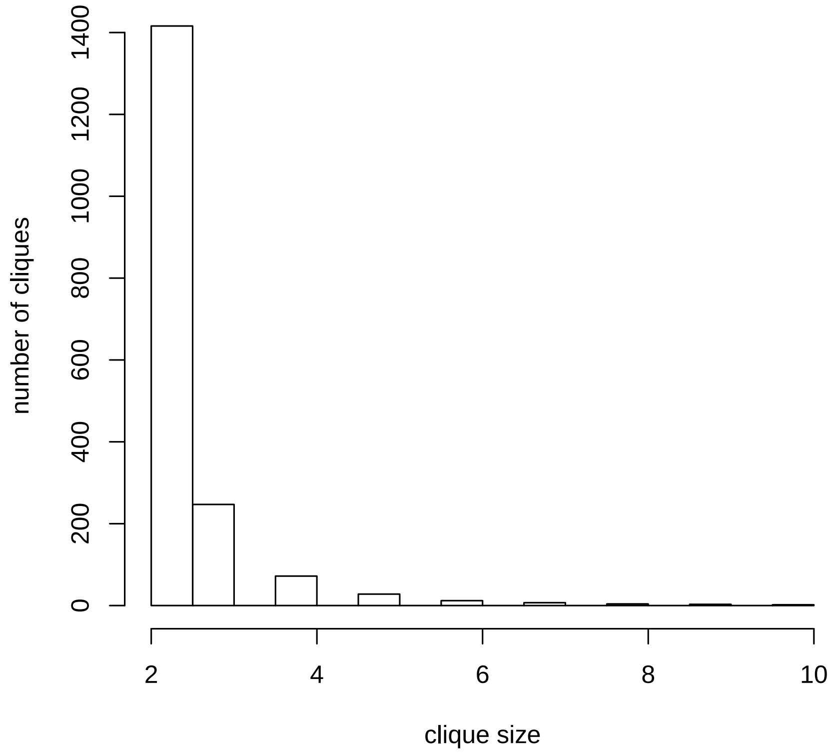

on a machine with 8 dual-core AMD Opteron 8218 (2.66 GHz) CPUs, and 128GB of RAM.

Table 1: Replanning modules

module	probability
best past plan	0.8
re routing	0.1
durations and mode optimisation	0.1

Table 2: Running time statistics

	without selection
$n_{threads}$	15
n_{iter}	60
total runing time	42 : 31 : 07
avg. replanning time per iteration	00 : 27 : 38

4.3 Output

4.3.1 Influence of constraints on the score of joint plans

When agents travel together, two constraints on the plan structure may have a major influence on the score of resulting plans, namely synchronisation and mode chaining constraints (a passenger will typically be unable to travel with his personal car in a (sub-)tour in which he is passenger).

Using a micro-simulation approach allows to simulate such constraints pretty easily. One may however wonder in which proportion such constraints influence the results. To do so, we ran 3 simulations on the same scenario, with the following settings:

1. passenger and driver plans are not synchronised, and passengers may use car in passenger subtours. This is somehow similar to the previous MATSim approach, were a fixed fraction of agents traveled with "ride" mode on whole subtours, experiencing car travel times without influencing congestion (Horni *et al.* (2011)).
2. passenger and driver plans are not synchronised, but passengers cannot use car in the subtours were they are passengers.
3. plans are synchronised, and mode constraints apply.

No waiting is actually simulated, so that passengers in unsynchronized plans will be moved even if the driver is not here.

Figures 2-3 illustrate the influence of plan synchronisation.

Figures 4-5 illustrate the influence of mode chaining constraints.

Figure 6 shows the scores of the plans with passenger trips at equilibrium. Mode chaining constraints seem to have a strong impact on the quality of a passenger plan: the impossibility to

Figure 2: Plan of clique 35685 (unsynchronised with mode constraints)

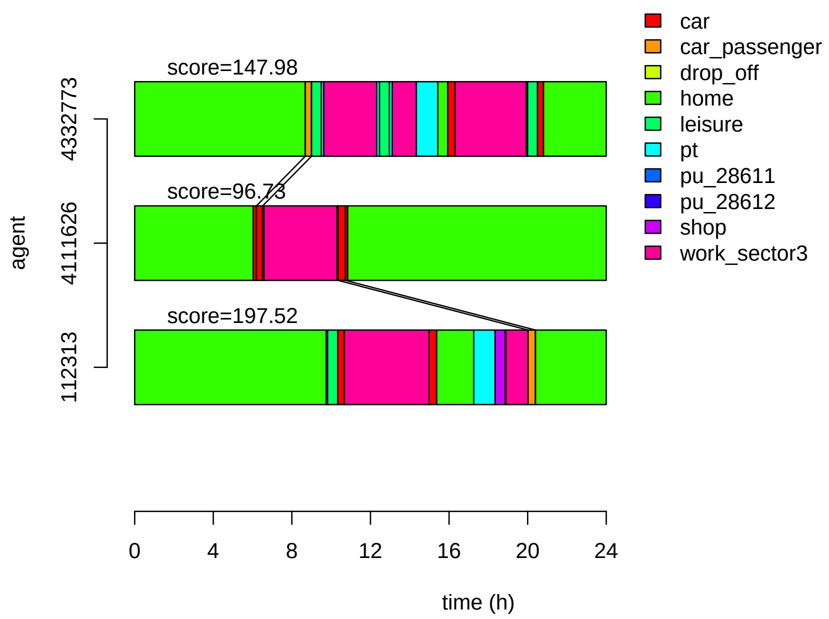

Figure 3: Plan of clique 35685 (synchronised with mode constraints)

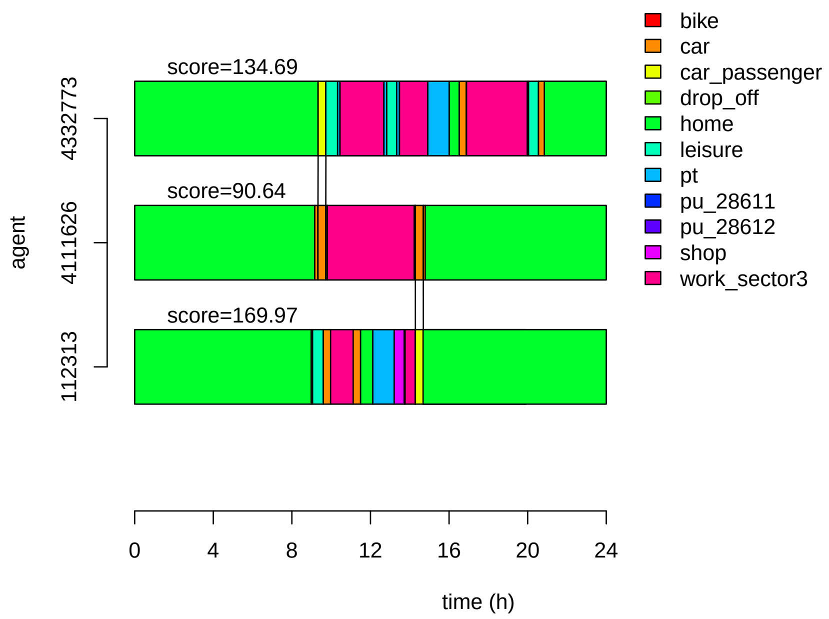

Figure 4: Plan of clique 482747 (unsynchronised without mode constraints)

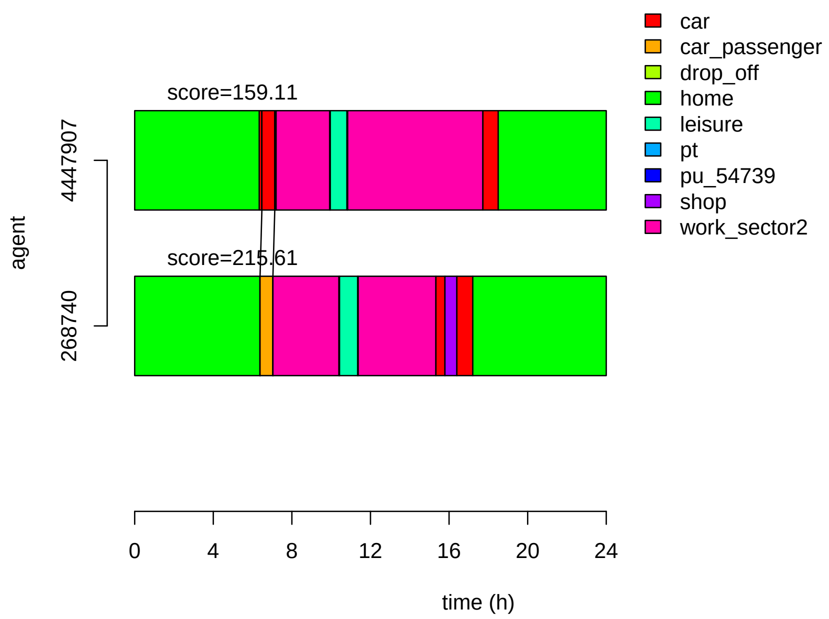

Figure 5: Plan of clique 482747 (unsynchronised with mode constraints)

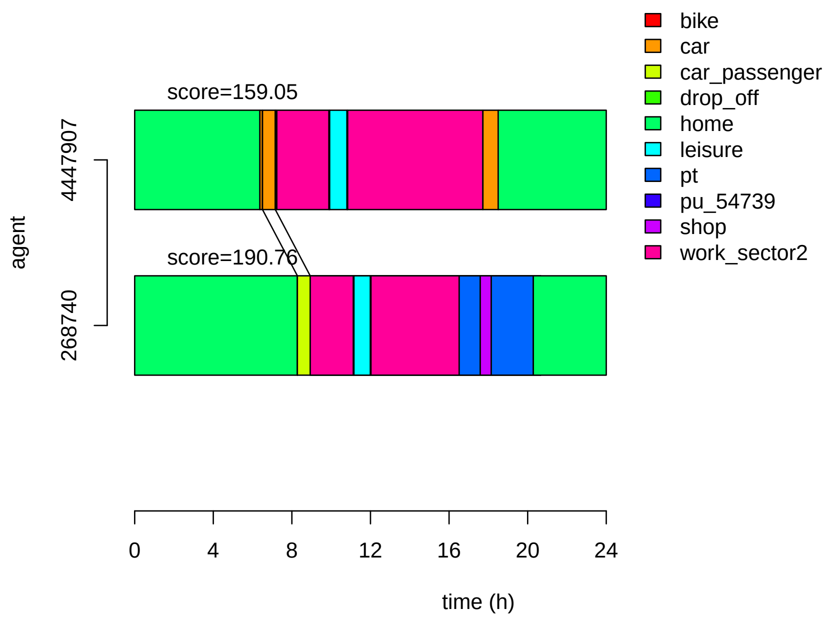

Figure 6: Scores of plans with passenger trips

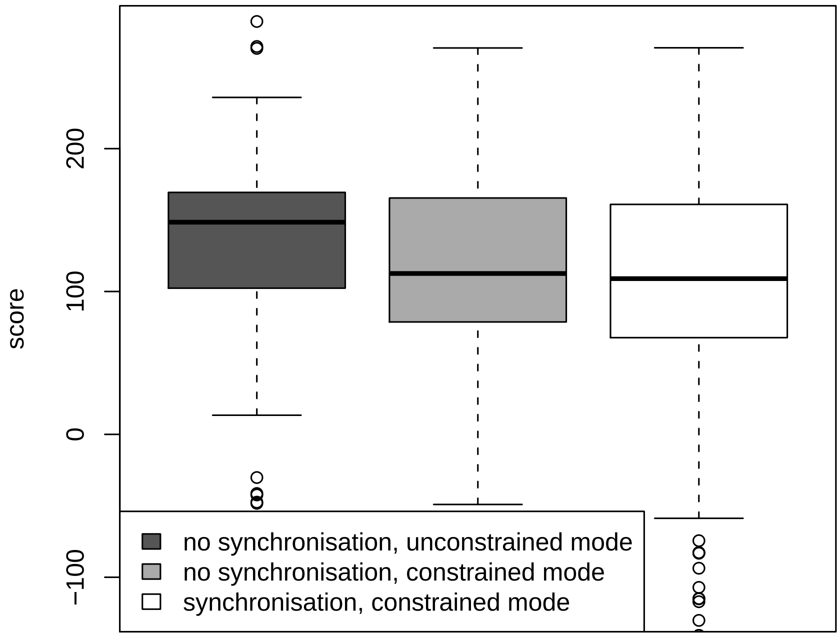

take the car afterwards is thus an important factor to take into account.

Synchronisation between co-travelers seems to have a much minor influence here. However, we remind that passenger and drivers were matched based on their optimal plans in a run without joint trips, so that only trips sufficiently close in time and space were grouped to create joint trips.

4.3.2 Influence of joint trips on scores and travel times

The following section focuses on scores and travel time with and without joint trips, at the steady state of the MATSim process. However, due to the stochastic nature of the relaxation process, analysing the differences in the plans of a particular agent between two runs would not make sense. Thus, the improvements induced by joint trips are computed in the following way: given the observed travel times in the last iteration:

- activity durations and modes for the selected plans at the steady state are optimized

Figure 7: Travel time improvements induced by passenger trips

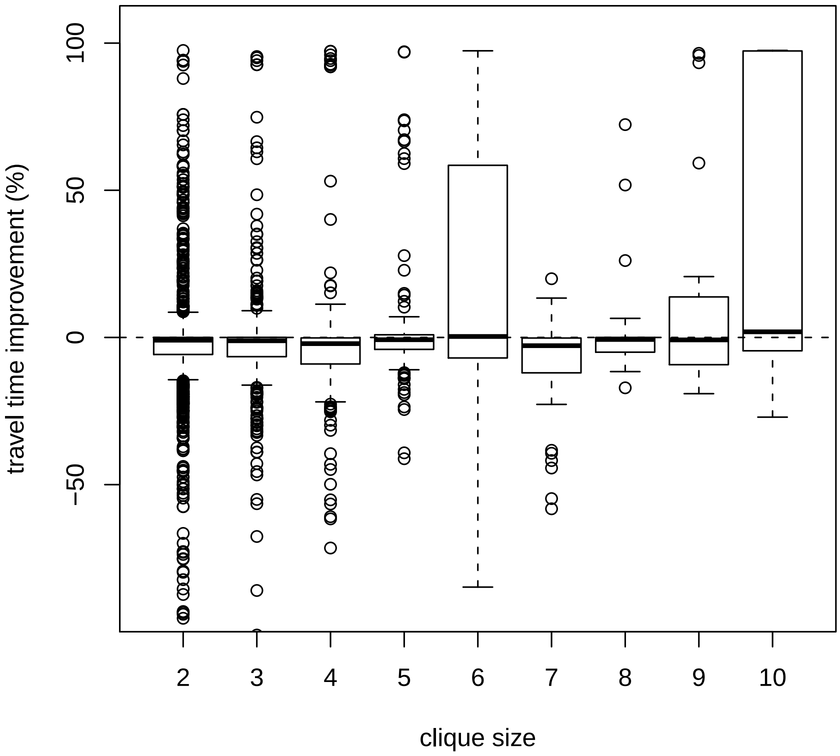

- for plans with joint trips, joint trips are removed, and activity durations and modes are re-optimized

this way, we obtain two scored plan, corresponding to the optimal strategy of the agent, given the strategy of others, with and without joint trips. When not stated otherwise, the presented results correspond to synchronised plans with mode constraints.

Figure 7 presents the variation of the travel times induced by passenger trips, that is, the difference in travel time between a passenger trip and the second best mode. The median of all improvements is close to 0. Some improvements are actually negative. This is probably due to the fact that most agents have a car: being driven cannot really improve travel times, and switching to car actually allows to choose a possibly less congested route. Good travel time improvements may correspond to agents without a car available.

Figures 8-9 show the improvements in plan score, depending on whether mode chaining constraints are taken into account or not. It is interesting to note that without mode constraints, the effect of passenger travel on the score is quite low and follows the same observations as travel time. However, when taking into account mode constraints, being a passenger can actually make

Figure 8: Score improvements induced by passenger trips, with and without mode chain constraints

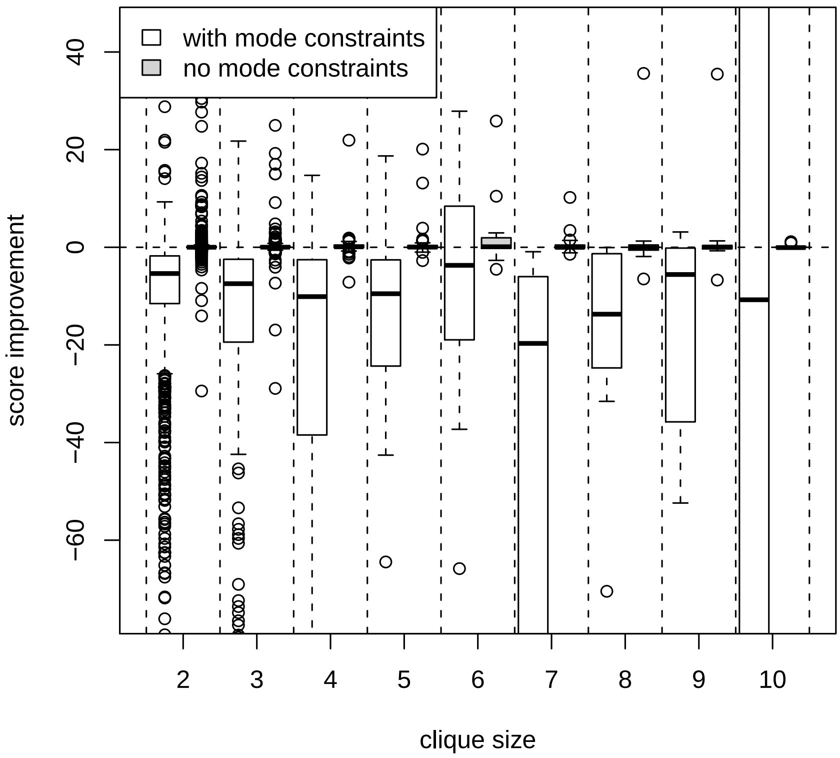

the score much worse. This should not be a surprise, but it points both the necessity to take such constraints into account, and to include other variables in the utility to actually reproduce the observed willingness to travel together (which may come from social closeness, willingness to reduce the cost of travel, limited number of cars in an household. . .).

Figure 9: Score improvements induced by passenger trips, with and without mode chain constraints (Zoom)

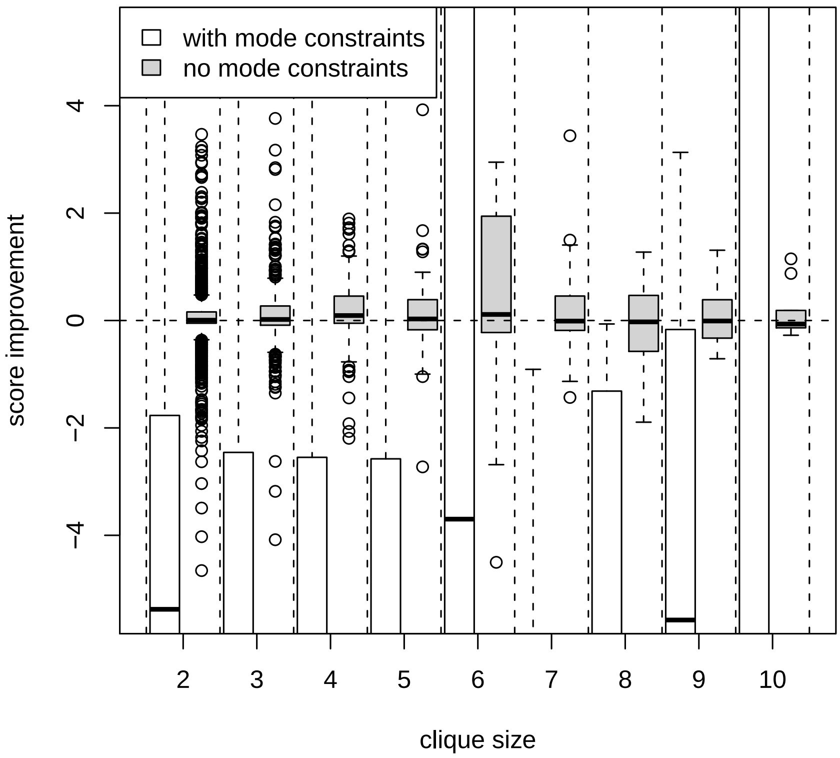

5 Conclusion and further steps

In this paper, we presented an approach to simulate joint trip in a multi-agent transport simulation, as well as an implementation of this approach within the MATSim software. This approach is a first step in considering collaboration in a user-equilibrium framework.

The output of a test run on a medium-scale scenario show that the implementation is able to produce meaningful results. It also allowed to reveal some interesting properties of joint travel in microsimulation. The more prominent is the influence of mode chaining constraints: our runs showed that considering the fact that being passenger prevents from using the car for the rest of the tour has a dramatic influence on the scores of passenger plans, and thus on the attractiveness of joint travel. It also incites to consider other variables than travel time when simulating joint trips, such as fuel costs or willingness to share time with social contacts.

Those remarks open a wide range of improvements to make to the model.

First, the utility function for daily plans must be modified to consider monetary costs or social relationships. Thus, a crucial future step will be to include this kind of dimension in the utility function, and to calibrate this utility.

Second, the implementation only works on pre-defined joint trips. It is planned to search for an efficient way to generate pertinent joint trips during the iterations.

Moreover, the current approach is to consider joint trips in pre-identified groups, to which we referred as cliques. A central, and strong, assumption in our approach is that "utility transfers" exist between clique members, making the existence of a clique-level utility function possible. This assumption was included in the model in order to take into account the high dependency of the scores of a plan on the actual co-travelers plans. However, it is planned to relax this assumption, passing again to individual utility maximisation, but using special plan selection modules, where the choices of plans by co-travelers are correlated. Utility transfers could still be represented explicitly, whenever the assumption is applicable.

Finally, no validation against measured data was undertaken yet. Once the approach is able to actually generate joint trips, a validation against data on intra-household ride sharing statistics should be undertaken.

6 References

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.

Glover, F. (1989) Tabu search - part 1, *ORSA Journal of Computing*, **1** (3) 190–206.

Glover, F. and E. Taillard (1993) A user's guide to tabu search, *Annals of Operations Research*, **41** (1) 1–28.

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.

Horni, A., B. J. Vitins and K. W. Axhausen (2011) The Zurich scenario: A technical overview, *Working Paper*, **687**, IVT, ETH Zurich, Zurich.

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

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.

Roughgarden, T. and É. Tardos (2002) How bad is selfish routing?, *Journal of the ACM*, **49** (2) 236–259, March 2002.

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.