High-Resolution Destination Choice in Agent-Based Demand Models

Date of submission: 2011-11-15

Andreas Horni
Institute for Transport Planning and Systems (IVT), ETH Zürich, CH-8093 Zürich
phone: +41-44-633 31 51
fax: +41-44-633 10 57
horni@ivt.baug.ethz.ch

Kai Nagel
Transport Systems Planning and Transport Telematics, TU Berlin, D-10587 Berlin
phone: +49-(030) 314 23308
fax: +49-(030) 314 26269
nagel@vsp.tu-berlin.de

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

Words: 6226 words + 4 figures = 7226 word equivalents

ABSTRACT

This paper describes the modeling of destination choice for discretionary activities in a multi-agent transport simulation, using MATSim as an example. MATSim is based on utility maximization. Randomness was included implicitly and in an uncontrolled way through the stochasticity of the simulation process, and sometimes through a logit choice model. Unobserved heterogeneity is now added directly to the utility function through a random error term. Importantly, those random error terms are *quenched*, i.e., they will be the same for repeated executions of the choice model. Real-world simulation experiments for Zurich show that this substantially improves results.

High-resolution destination choice for large-scale microsimulations raises several technical issues; pragmatic engineering solutions have been developed or applied to cope with them. These solutions are described in technical detail to assist in the further development of similar microsimulations.

INTRODUCTION, PROBLEM AND GOAL

Computational Process vs. Utility-Based Models

There are two general transport microsimulation types: computational process models and utility-based models. Utility-based models focus on the definition of final decision *outcomes*. These outcomes maximize the decision maker's utility, subject to constraints. Computational process models, on the other hand, assume that decisions are captured realistically by focusing on the decision *process*, guided by heuristic decision rules.

Destination choice in utility-based transport microsimulators is mainly based on discrete choice models, where the choice set is constructed obeying various constraints, such as Hägerstrand's space-time prisms (see e.g., *(1, 2, 3, 4)*).

MATSim *(5)* belongs to the strand of utility-based transport simulation frameworks. The persons' activity chains are iteratively simulated, adapted and evaluated by a utility function as described later.

Heterogeneity

Transport models need to adequately treat heterogeneity in the context of travel decisions, usually modeled by random error terms for every person-alternative pair and by probability distributions for the model coefficients (mixed models).

This is explained in more detail by a prototypical example: a toll road, with a non-toll, but slower, alternative. In the *absence of congestion*, the Nash equilibrium solution has, for any given origin-destination (OD) pair, either all demand on the toll road, or all demand on the non-toll road. This is not realistic; for certain values of toll and travel time, the traffic streams will split onto both options.

A typical approach to the problem assumes person-specific unobserved attributes for every alternative modifying the utility; i.e., one has:

$$U_{pi} = V_{pi} + \epsilon_{pi} \,, \tag{1}$$

where p is the person index, and i the index for the alternative. As usual, V denotes the systematic part of the utility (the same for every person of a given OD relation), ϵ is the random offset, and U is the resulting utility on which the user equilibrium will be based. As is well known, one typically progresses assuming that the ϵ_{pi} are independently and identically Gumbel distributed, leading to the logit choice model, and from there, to the stochastic user equilibrium (SUE).

An alternative is to assume that there are person-specific coefficients that modify the decision. Let us, for example, take individual values-of-time VoT_p. For this, let us assume that V is of the form:

$$V_{pi} = -\alpha \, t_{pi} - \gamma_p c_{pi} \,,$$

where t is travel time, c the toll ("cost"), α the weighting factor of time, and γ_p the weighting factor of the toll. Importantly, γ_p depends on the person. Assuming that travel times and toll payments are the same for everybody, then alternative i is better than alternative j if

$$V_{pi} > V_{pj} \,,$$

$$-\frac{\alpha}{\gamma_p} t_i - c_i > -\frac{\alpha}{\gamma_p} t_j - c_j \,,$$
$$\frac{c_i - c_j}{t_j - t_i} < \frac{\alpha}{\gamma_p} =: VoT_p \,,$$

where the last equation is only valid when $t_j > t_i$. Assuming then that i is the toll option, one finds that i is accepted only when the payment divided by the time gain is less than the person-specific value of time. Persons with a large value of time will tend to use the toll road; persons with a small value of time will tend to use the non-toll, slower road.

However, it is *not* possible to use the taste variation approach *without* the unobserved attributes approach, since not all heterogeneity in behavior can be explained by taste variations. Drawing on the example above, there may be an OD relation where the toll road is both slower and more expensive, i.e., it is pareto-dominated. That some users may use it anyway can thus, within the utility-maximizing framework, only be explained by additional, unobserved, attributes.

Repeated Draws: Quenched vs. Annealed Randomness

Including Randomness in the Microsimulation

One assumed advantage of microsimulation is the conceptually straightforward inclusion of heterogeneity. In the first instance, one can, whenever it is needed, either

- randomly draw from a choice model given as probability distribution or
- randomly generate an ϵ_{pi} for every person-alternative pair and select an alternative i as $\underset{i \in choice\ set}{\operatorname{argmax}} U_{pi}$.

For mixed logit models, one can also draw a VoT_p from a distribution *(*e.g., *6, 7)*.

In both cases, however, problems with repeated draws must be solved. Repeated draws mean that the same individual p is repeatedly faced with an identical choice, a frequent situation in iterative models. Obviously, the VoT_p of individual p should not change from one such draw to the next, and similarly, the ϵ_{pi} should remain fixed once they have been drawn for the first time. For the same reason, neither the ϵ_{pi} nor the VoT_p should change during the introduction of a policy measure, except when it affects them directly.

In physics, this would be called "quenched" randomness; all randomness is computed initially and then attached to particles or locations, rather than instantaneously generating it, which would be called "annealed" randomness.

Implementing Quenched Randomness

Quenched randomness can be achieved by applying one of the following two strategies:

- (a) Freezing the applied *global* sequence of random numbers, meaning that a Monte Carlo method with the same random seed is used before and after the introduction of a policy measure and over the course of iterations. Thus, the VoT_p and ϵ_{pi} should come out the same way *before* and *after* the introduction of the policy measure. Differences in the outcome can thus be directly attributed to the policy measure.
- (b) Computing and storing a separate ϵ_{pi} for every combination of person p and alternative i.

The authors reviewed relevant literature, but could not determine strategies applied in each case in other large-scale transport microsimulations. Through personal e-mail communication with the simulator authors, some answers emerged: in AMOS and OpenAMOS *(8, 9)* (a) is applied. In Albatross *(10)* both (a) and (b) have been applied. For the NYC activity-based microsimulation *(11)* in most cases (a) is used, although they recently switched to (b). The Tel Aviv model *(12)* is based on (a). The Sacramento and Portland models developed by Mark Bradley and John L. Bowman besides others *(*e.g., *13, 14)* apply (a).

Both strategies have flaws. Approach (a) is only an option if one is certain about all aspects of the computational code. Importantly, one additional random number, drawn in one run but not in the other, completely destroys the "quench" for all decisions computed later in the program.

Thus, approach (b) could be more robust in practice. However, for large numbers of decision makers and/or alternatives, storing error terms is difficult. For destination choice, one quickly has 10^6 decision makers and 10^6 alternatives, resulting in 4×10^{12}Byte = 4TByte of storage space.

One may argue that this should not be a problem, since a normal person will rarely consider more than the order of a hundred alternatives in their choice set, reducing the computational problem. Aside from the necessity of storing every decision maker's choice set, this converts the computational problem into a conceptual one, since a good method to generate choice sets then needs to be found. With more conceptual progress, this may eventually be an option, but at this point, a conceptually simpler approach is preferred.

As far as the authors know, this set of problems has not been discussed in existing literature. In this paper, we present a relevant solution for the computational problem associated with approach (b).

MATSim and Heterogeneity

The utility function used until now (for route and time choice) did not contain a random error term. However, a certain amount of randomness (i.e., unobserved heterogeneity) implicitly entered the model: Two identical persons with the same origin and destination may still end up with different routes according to the random order in which they undergo the replanning. Essentially, this means that a random term is added implicitly to the choices. Also, some investigations applying MATSim use a logit choice model, thus adding another element of randomness.

However, this randomness is introduced in an unsystematic manner. Future investigations are thus necessary for these choice dimensions, but this paper is focused on destination choice.

Heterogeneity in the New Destination Choice Model

The newly added MATSim destination choice model, like most operational destination choice models, does not take locational competition into account. Thus, constraints coupled to competition cannot introduce unobserved heterogeneity the way they can for time and route choice.

This means that every time a synthetic traveler is up for leisure destination choice (for example), he or she will switch, if available, to a new leisure location generating a higher score. Without additional measures, such a leisure location will almost always either be closer to home, or closer to a trip that is done anyway, such as work-to-home. With this setup, leisure trip distances keep decreasing over the iterations.[1] In practice, not taking into account the unobserved heterogeneity leads to a dramatic underestimation of total travel demand (visible in too-short travel distances, too-short travel times and underestimated link volumes as compared to count data) as shown below and in *(16)*.

The intuitive explanation is that synthetic travelers do not differentiate between possible destinations; one leisure or shopping location is as good as any other. An "observed" way to address this problem might be to differentiate these facilities by type, e.g., leisure facilities into tennis, swimming, rock climbing, parks, etc. However, in many cases, neither the land use nor the observed tastes are available, and *even then* some unobserved heterogeneity will remain.

To conclude, when including destination choice, explicit random error terms should be used to be compatible with econometric discrete choice methodology *(17)*.

This Paper's Goal

This paper will introduce an operational destination choice approach to overcome the problems identified below for high-resolution microsimulations. To test the new approach, a suitable module is implemented in MATSim.

In this paper, unobserved heterogeneity is explicitly incorporated as an error term in the utility function. Random taste parameters are not discussed, but are left for future work.

The following technical issues are raised:

First, as discussed above, with MATSim's iterative structure, drawing from error distributions is not straight-forward. This holds for all iterative procedures. If the error term is drawn *per iteration* with an arbitrary random seed, there is no convergence toward a stable solution. An efficient mechanism to assign a fixed individual error term per person-alternative pair ϵ_{pi} (for person p and destination i) will be presented.

Second, travel times are a very important determinant in destination choice. Calculating travel times includes routing, which is computationally very expensive. This makes computation of travel times for many alternatives, as it is necessary for destination choice, difficult, or even unfeasible. This problem is investigated and a first solution is presented.

METHOD: DESIGNING THE MATSIM DESTINATION CHOICE MODULE

MATSim

MATSim is an activity-based, extendable, open source, multi-agent simulation toolkit implemented in JAVA and designed for large-scale scenarios and is a co-evolutionary model. In competition for space-time slots on transportation infrastructure with all other agents, every agent iteratively optimizes its daily activity chain by *trial and error*. Every agent possesses a fixed amount of day plans memory, where each plan is composed of a daily activity chain and an associated utility value (in MATSim, called *plan score*). Computation of plan score is compatible with micro-economic foundations and is described in more detail below.

Before plans are executed on the infrastructure in the network loading simulation *(*e. g., *18)*, a certain share of agents (usually 10%) is allowed to select and clone a plan (here: the plan with the highest score per agent) and to subsequently modify this cloned plan. Three

[1] Note that this occurs only when iterative choice dimensions include destination choice. There are very few MATSim papers doing this *(15, 16)*.

choice dimensions are considered: time choice *(19)*, route choice *(20)*, and destination choice as described in this paper.

If an agent ends up with too many plans (here set to “5 plans per agent”), the plan with the lowest score (configurable) is removed from the agent’s memory. One iteration is completed by evaluating the agent’s day described by the selected day plans.

If an agent has obtained a new plan, as described above, then that plan is selected for execution in the subsequent network loading. If the agent has *not* obtained a new plan, then the agent selects from existing plans. The selection model is configurable. In many MATSim investigations, a model generating a logit distribution is used. However, for this paper, agents will select the plan with the highest score.

The MATSim Utility Function

The basic MATSim utility function was formulated in *(21)*, from the *Vickrey* model for road congestion as described in *(22)*. Originally, this formulation was constructed for departure time choice. Several studies e.g., *(23)*, indicate that the extended function is productive for modeling time choice of complete days, including route choice. It has thus been adopted as the starting point for handling destination choices.

The utility of a plan U_{plan} (described in detail in *21)* is computed as the sum of all activity utilities $U_{act,q}$ plus the sum of all travel (dis)utilities $U_{trav,q}$

$$U_{plan} = \sum_{q=1}^{n} U_{act,q}(type_q, start_q, dur_q) + \sum_{q=2}^{n} U_{trav,q}(loc_{q-1}, loc_q)\,,$$

where $type_q$, $start_q$ and dur_q are the type, start time and duration of the activity q respectively. Utility of an activity is essentially dependent on activity duration. Details are given in *(23)*. The parameter setting described in this paper is applied to the real-world simulation configurations 2 and 3. For configurations 0 and 1, different utility functions are employed, described at the appropriate location.

Earlier MATSim Destination Choice Solutions

The Swiss Census of Population 2000 *(24)* can identify home and work locations for every Swiss resident at hectare and municipality level respectively. Clearly, such information can not be logged for discretionary activities. However, to run an activity-based simulation, reasonable destinations for these activities must be assigned. First, a simple neighborhood search, as described in *(23)*, was employed in a preprocessing step. That approach is not part of the optimization process and does not accurately model destination choice.

A first improvement in destination choice—including it in the optimization process—was introduced by *(15)*, based on Hägerstrand’s time geography. However, unobserved heterogeneity was not taken into account in that module or in MATSim. Thus, a significant underestimation of travel demand resulted and the module could not be productively employed. Furthermore, that module is based on local search. Local search applicability, however, is questionable on destination choice utility space. Its specific characteristics are explained later.

Incorporating Heterogeneity by Individual Error Terms in an Iterative Model

In iterative models, ensuring a stable choice for the same person doing the same choice over the course of the iterations can be achieved as follows. Fixed random error terms are assigned to every person-destination pair pi. These terms can be randomly assigned in a preprocessing step and held constant over the course of iterations. The optimization is then performed as a deterministic search based on the resulting utility function. In fact, this can be seen as a return to the roots of random utility modeling – rather than absorbing the ϵ_{pi} into the choice model, they are now explicitly generated.

However, when trying to store these error terms directly, an infeasible storage effort results for large-scale scenarios as shown earlier. Instead, the same *stable* error term can be *re*-calculated on the fly by using the random seed $s_{pi} = g(k_p, k_i)$. The distribution of these seeds is essentially irrelevant. In this paper k_p is a fixed uniformly distributed value per person p, and k_i is a fixed uniformly distributed value per destination i. In this work $g(k_p, k_i) = (k_p \times k_i) \times v_{max}$ is used. v_{max} is the maximum (long) number that can be handled by the specific machine.

To evaluate utility for a person p visiting the destination i a sequence of Gumbel-distributed random numbers seq_{pi} is generated on the fly for every person-alternative pair using the seed s_{pi}. The error term ϵ_{pi} is then derived from the m^{th} element of the sequence $seq_{pi}[m]$. Here, m is set to 10. This procedure is valid as the set of all m^{th} elements of all different sequences is also a pseudo-random sequence following the same distribution as the sequences seq_{pi}. In this work, a standard Gumbel distribution (i.e., location parameter $\mu = 0$ and scale $\beta = 1$) is applied. It is first scaled to produce a 1.0 standard deviation. Second, as no utility function estimation is yet available, calibration of error terms is performed, where two parameters $f_{shopping}$ (here set to 0.95) and $f_{leisure}$ (here set to 1.35) are used. Clearly, *true* random number generators relying on physical phenomena, such as hardware temperature, are not applicable.

Designing the Destination Choice Module

Essentially, MATSim is based on random mutation. However, the huge number of available alternatives for all choice dimensions (and hence curse of dimensionality) makes the introduction of optimizing mechanisms indispensable, i.e., mechanisms that return a "good" rather than a random alternative from the set of available alternatives.

Search Space and Search Method

The discrete search landscape is characterized by random noise because error terms are not (or only locally) spatially correlated (see Figure 1). For such problems, efficient search methods, such as local search methods, generally do *not* work. Furthermore, in the model and, in reality, utility contributed by the error term is *unlimited*. The search space for potential destinations is hence *unlimited*. Unfortunately, exhaustive search usually produces prohibitively large computation efforts for large-scale scenarios. Thus, the application of problem-tailored heuristics and approximations is unavoidable.

A first attempt to narrow down individual search space Γ_{pq} for person p and activity q is as follows. In discrete choice theory, individual p chooses alternative i, producing maximum utility for activity q:

$$U_{iqp} \geq U_{jqp}, \forall j \in choice\ set\,,$$

that is

$$V_{iqp} + \epsilon_{iqp} \geq V_{jqp} + \epsilon_{jqp}, \forall j \in choice\ set\,,$$

where V denotes the deterministic part of the utility function. V is usually composed of travel effort V_{travel} and utility for performing an activity $V_{perform}$. Hence, the above formula gives:

$$V_{travel,iqp} + V_{perform,iqp} + \epsilon_{iqp} \geq V_{perform,jqp} + V_{travel,jqp} + \epsilon_{jqp}, \forall j \in choice\ set\,.$$

In MATSim, $V_{perform,jqp}$ is equal for all destinations j if the performed activity time is equal and it decreases with increasing travel effort. Hence, maximum potential travel effort is equal to

$$\epsilon_{pq,max} := \max_{\omega \in choice\ set} \epsilon_{\omega pq}\,.$$

This defines the upper-most boundary of the search space. Actually, search space is restrained by the travel effort to reach the *destination* associated with the largest error term. However, it is assumed that this stopping condition is seldom satisfied, and thus, its efficiency does probably not justify the additional implementation and computation complexity. Instead, the search space can be further restrained under the natural assumption that an activity is dropped if it does not generate positive utility at least for one destination, i.e.,

$$V_{travel,jpq} + V_{perform,jpq} + \epsilon_{jpq} \stackrel{!}{>} 0\,.$$

The dependency of $V_{perform}$ on V_{travel} is difficult. To make things worse, $V_{perform}$ is usually non-linear. Fortunately, $V_{perform}$ can be omitted when searching for an *upper* bound for the accepted travel costs. Clearly, $V_{perform}$ is larger for closer locations: The longer the trip takes, the less time there is to perform the activity. In other words, the benefit decreases by traveling due to travel costs V_{travel} and opportunity costs (loss of $V_{perform}$). This loss must be at least compensated by the error term for a person to choose a more distant destination and to not stay at the current (closer) location: A person only travels farther if that effort produces a net benefit. An *upper bound* for maximum search space can thus be found by considering only the compensation of the travel costs, i.e., by ignoring the opportunity costs (lost activity performing time).

Hence the above equation becomes after setting $\epsilon_{jpq} := \epsilon_{pq,max}$ and rearranging:

$$V_{travel,pq} > -\epsilon_{pq,max}\,.$$

Note that V_{travel} is negative. The destination index j can be omitted, as the formula is not destination-dependent.

Now assume:

$$V_{travel,pq} = \beta_{pq} \times f_{pq}(t, d, m)\,,$$

where β_{pq} is the individual cost coefficient for person p and usually negative. $f_{pq}(t, d, m)$ is the travel cost function, usually composed of time (t), distance (d) and monetary (m) costs.

Maximum travel costs defining search space are thus given after rearranging by:

$$f_{pq}(t, d, m) < \frac{-\epsilon_{pq,max}}{\beta_{pq}}\,.$$

The inequality sign changes as β_{pq} is negative. In this paper, linear travel distances $V_{travel,pq} = \beta_{distance,pq} \times distance$ are used as travel disutilities.

Therefore, the above equation translates to:

$$distance^{max}_{pq} < \frac{-\epsilon_{pq,max}\ [\ ]}{\beta_{distance,pq}\ [\frac{1}{m}]}\,. \tag{2}$$

This approach is promising, as very large values for $\epsilon_{pq,max}$ are rare (Figure 2(a)), meaning that a huge space must be searched for only a few persons. The search space Γ_{pq} is constructed as follows. Let us assume that for the activity q of person p, a new location l_{pq} has to be found. Γ_{pq} can then be defined as a circle whose center is the mid-point between the preceding activity l_{pq-1} and the succeeding activity l_{pq+1}.

The radius of the circle is set to:

$$r_{\Gamma_{pq}} = (distance(l_{pq-1}, l_{pq+1}) + distance^{max}_{pq})/\psi\,.$$

The most productive value for ψ is not yet apparent. For every discretionary trip, there is a $distance^{max}_{pq}$ that person p is willing to travel at most. Looking at an individual discretionary tour with fixed and identical locations $l_{pq-1} = l_{pq+1}$ clearly $distance^{max}_{pq}$ includes the return trip and $\psi = 2$ is thus a natural choice. But, for consecutive multiple discretionary activities the search space is probably larger, and ψ is thus smaller. However, essentially, the value of ψ is subject to calibration and needs further research. In this paper, $\psi = 2$ is used.

It is crucial that Γ_{pq} can be computed fast and that all destinations actually accessible are contained. On the other hand, only computation times, but not the quality of the results, are influenced if destinations that are actually inaccessible are included in the evaluation. For that reason, it is possible to approximate the travel distance $distance^{max}_{pq}$ by the straight-line distance. This distance can then be computed once in a preprocessing step.

Convergence Speed, MATSim Best Response and Computation Times

Normally, exactly one alternative per choice dimension is evaluated in each MATSim iteration. This procedure is reasonable as the persons interact in the infrastructure and influence each others' choices, which can be regarded as a feedback mechanism.

For destination choice, due to numerous available destinations in the search space, a huge number of iterations is required, resulting in a very low *convergence speed.*

As long as the change of travel costs between succeeding iterations is not too large, multiple search space destinations can be evaluated per person and per iteration. Normally, the relatively small share of agents who re-plan, keep the inter-iteration changes small. Thus, increasing the number of evaluated alternatives per iteration might be feasible. This reduces the number of iterations and substantial costs associated with simulation of network loading.

Despite this simulation time reduction, computation times are still infeasible and further speed-ups are necessary. For the 10% Zurich scenario with approx. 68'000 persons, one iteration takes roughly 20 hours, even when using multiple processor cores. Most computation time is due to calculation of travel times, i.e., due to routing, in the context of large alternatives sets. To reduce these huge routing costs, the following procedure is applied.

Let us assume that location l_q of activity q is changed, where all other plan activities are fixed. Travel times for routes between activity location l_{q-1} and all potential locations l_q can be exactly and efficiently computed by Dijkstra's algorithm because it efficiently computes the best routes from one location to *all* other locations in the network. Travel times of the best routes between activity locations l_q and l_{q+1} are computed by running Dijkstra's algorithm *backwards*, using an average estimated arrival time as initial time. This is an approximation, as the arrival time at l_{q+1} is different for different locations l_q.

To reduce possible approximation errors, a *probabilistic* best response is applied. Search space destinations are evaluated as described above; then a random choice weighted by these approximated scores is performed. The plan containing the new choices is finally simulated and eventually scored, based on *exact* travel times by the MATSim iteration scoring. This approach is justified by the assumption that, during the course of the iterations, the probabilistic choice probably reduces—or even compensates for—the errors incurred by approximating travel times as described above.

However, the probabilistic choice brings back the problem of slow convergence. If every alternative in the search space is chosen with probability greater than zero, this huge set necessitates a large number of iterations. For reasonable convergence, the probabilistic choice must be performed on a reduced choice set. Thereby, restraining the choice set to the ϕ destinations producing the highest approximate plan scores is natural. ϕ is essentially dependent on the approximation error done by estimating travel times. For our purposes ϕ is set to 30 but this value also needs further research.

With this procedure the required computational effort is dramatically reduced, allowing application of destination choice to large-scale scenarios. One iteration of the 10% Zurich scenario takes roughly 25 minutes (instead of 20 hours). The simulation is run with 10 parallel JAVA threads and approximately 15GB of RAM. The Linux server is equipped with an Intel Xeon(R) processor, 3.33GHz, with 24 cores and 96GB of RAM.

RESULTS

The paper investigates the direct inclusion into the utility function of heterogeneity that was so far not present in MATSim. Solutions are presented to overcome serious computational difficulties in large-scale scenarios using examples integrated in both a small-scale scenario and a real-world scenario (the often used Zurich scenario as described below).

Synthetic Small-Scale Scenario

The synthetic scenario is composed as follows. The study area is a 20 km square. The network consists of 6561 nodes and 25920 directed links, building a grid of 6400 squares. The link capacities are 400 vehicles per hour. The scenario consists of 2000 persons, whose home locations are located at 40 central locations in the study area. 12960 shopping destinations are equally distributed over the study area, so that every square is connected with 4 destinations. Persons' activity chains contain two home activities with an intermediate shopping activity. During the iterations, time, route and destination choices are made. Initially, all shopping activities are performed at the home location. The initial end time of the first home activity is 11 o'clock and the desired shopping duration is arbitrarily set to 180 minutes. The opening times window is set narrow (from 9:30 to 14:30) so that persons cannot completely circumvent traffic jams. These values roughly represent Saturday shopping trips. Real data is applied in the Zurich scenario as described below.

The following two configurations are simulated:

- **Configuration 0:** $U = \epsilon$.
  The utility function applied for this configuration is made up only of error terms.
- **Configuration 1:** $U = \beta_{distance} \times distance + \epsilon$, with $\beta_{distance} = const = -0.0005[\frac{1}{m}]$. In this configuration a distance-based utility function is applied; i.e., a linear disutility distance term is added to the setting of configuration 0.

As travel times are not part of the utility function for the configurations 0 and 1, the travel time approximation described earlier is not applied for these configurations. The error terms ϵ are independently and identically (i.i.d.) Gumbel distributed with $\sigma_\epsilon = 1$.

Configuration 0

In this configuration choices are independent of travel costs. The resulting distance distribution is shown in Figure 2(b). Starting from the center, and increasing the travel distance, the number of destinations increases linearly. This means that also the probability for finding the destination generating maximum utility increases linearly with distance. As soon as the boundaries of the study area are reached, the probability of finding an even better alternative by increasing the travel distance again decreases. The relatively slow decrease is due to the geometrical setting. The area of potentially better destinations is a circle, where the study area is a square. Thus, some agents find their best option in the corner areas of this square, lying already outside of the circle mentioned before.

The distribution of the maximum score (utility) per person again follows a Gumbel distribution (see Figure 2(a)).

Configuration 1

In Figure 2(b), the distance distribution roughly follows a negative exponential distribution as observed in empirical data, such as the Swiss National Travel Survey *(25)*. This means that, in essence, realistic distance statistics can be produced very efficiently without imposing any complex boundary conditions *(*as in the earlier MATSim destination choice models, see *16)* or applying complex behavioral models.

When looking at real data, a relatively long distribution tail can usually be observed even for home-based round-trips; whereas for this scenario long distances are rare. Although the scenario is synthetic, this may be further evidence that the distance cost perception function might be better assumed to be non-linear with decreasing marginal costs for very long distances.

As mentioned earlier, this paper aims to provide a mechanism for sampling from discrete choice models in combination with optimization of complete day plans in a fully integrative way. A first verification step of this mechanism is provided here by re-estimating the utility function distance parameter $\beta_{distance}$. The choice set consists of all shopping destinations in the study area, with all 2000 agents included. Choices are given by the agents' destination choices for the relaxed state (final iteration 100). Consistency between the applied and the estimated coefficient can be observed in the significant results: $\hat{\beta}_{distance} = -0.000651$ with $\rho^2 = 0.285$.

Real-world Scenario: 10% Zurich Scenario

The 10% Zurich scenario is frequently used in MATSim development but also for projects in Swiss planning practice *(*e.g., *23, 15).* Simulation scenario demand is derived from the Swiss Census of Population 2000 *(24)* and the National Travel Survey for the years 2000 and 2005 *(25)*. A 10% sample of car traffic (including cross-border traffic) crossing the area delineated by a 30 km circle around the center of Zurich (Bellevue) is drawn, which results in almost 68'000 agents simulated. The activity location data set, comprising more than 10^6 home, work, education, shopping and leisure locations, is computed from the Swiss Census of Population 2000 and the Federal Enterprise Census 2001 *(26)*. The network from the Swiss National Transport Model *(27)* is used, consisting of 60'492 directed links and 24'180 nodes. A single day is simulated, with 3.35 average number of trips per agent. In total, 25'896 shopping activities and 40'971 leisure activities are performed. Comparable data is available in most countries from official sources, such as censuses, national travel diary studies and commercial sources, such as navigation network providers, yellow pages publishers or business directories.

The trip distance distributions are taken from the National Travel Survey for the year 2005 *(25)*, reporting 33'000 person days for Switzerland overall. Patterns for trips undertaken by persons visiting or living in the Zurich region compared to the complete set of trips are almost identical. However, due to the smaller sample size, the restrained set shows more noise. Thus, all trips in Switzerland are used. Additionally, traffic count data for 2004-2005 from automatic national, cantonal and municipal count data stations *(*e. g. *28)* are taken into account. Count data is evaluated in this paper for the area delineated by a circle with a 12km radius, containing 123 counted links.

The following two configurations are simulated:

- **Configuration 2:** $U = f(t_{activities}, t_{travel})$, i.e., excluding unobserved heterogeneity where $f(.,.)$ refers to the standard MATSim utility function described earlier.
- **Configuration 3:** $U = f(t_{activities}, t_{travel}) + \epsilon$, i.e., including unobserved heterogeneity.

Figures 3 and 4 show that both traffic counts and distances traveled for shopping and leisure trips are underestimated if error terms are excluded (configuration 2). This problem would intensify for weekend scenarios to be developed in the near future. Results indicate that incorporating error terms (configuration 3) is highly productive. A very good match between simulated and measured values can be achieved with only minimal calibration efforts. Similar results are achieved for link volumes compared to traffic count data. The median relative error of daily volumes (averaged over the 123 links) is reduced from almost -40% to approximately -20% (see Figure 4).

By applying the travel time approximation described earlier, the scenario is computable in reasonable time. It takes roughly 25 minutes per iteration, reaching a stable state after 100 iterations. However, the 10% Zurich scenario is still at the lower bound of typical MATSim projects. Thus, further speed-up mechanisms need to be researched.

CONCLUSIONS AND OUTLOOK

This paper introduces an operational destination choice approach to overcome problems identified for high-resolution microsimulations. To test the new approach, a suitable module is implemented in MATSim that may also be important for other iterative utility-maximizing transport simulations. With the direct inclusion of random error terms in the MATSim utility function some previously lacking unobserved behavioral heterogeneity is introduced.

As shown in a real-world scenario, destination choices can now be modeled realistically in MATSim, a crucial step for many transport research questions that must be investigated in the near future.

In the future, applied parameters need to be researched more comprehensively and more systematically calibrated. The utility function for shopping trips will be estimated in the context of *(29)* and the Zurich scenario will be enhanced by various destination attributes, reducing the error terms.

Other important topics include researching Monte Carlo sampling and its sampling errors associated with our stochastic simulation. The authors have begun to work on these issues *(30)*.

FIGURE 1 Search space

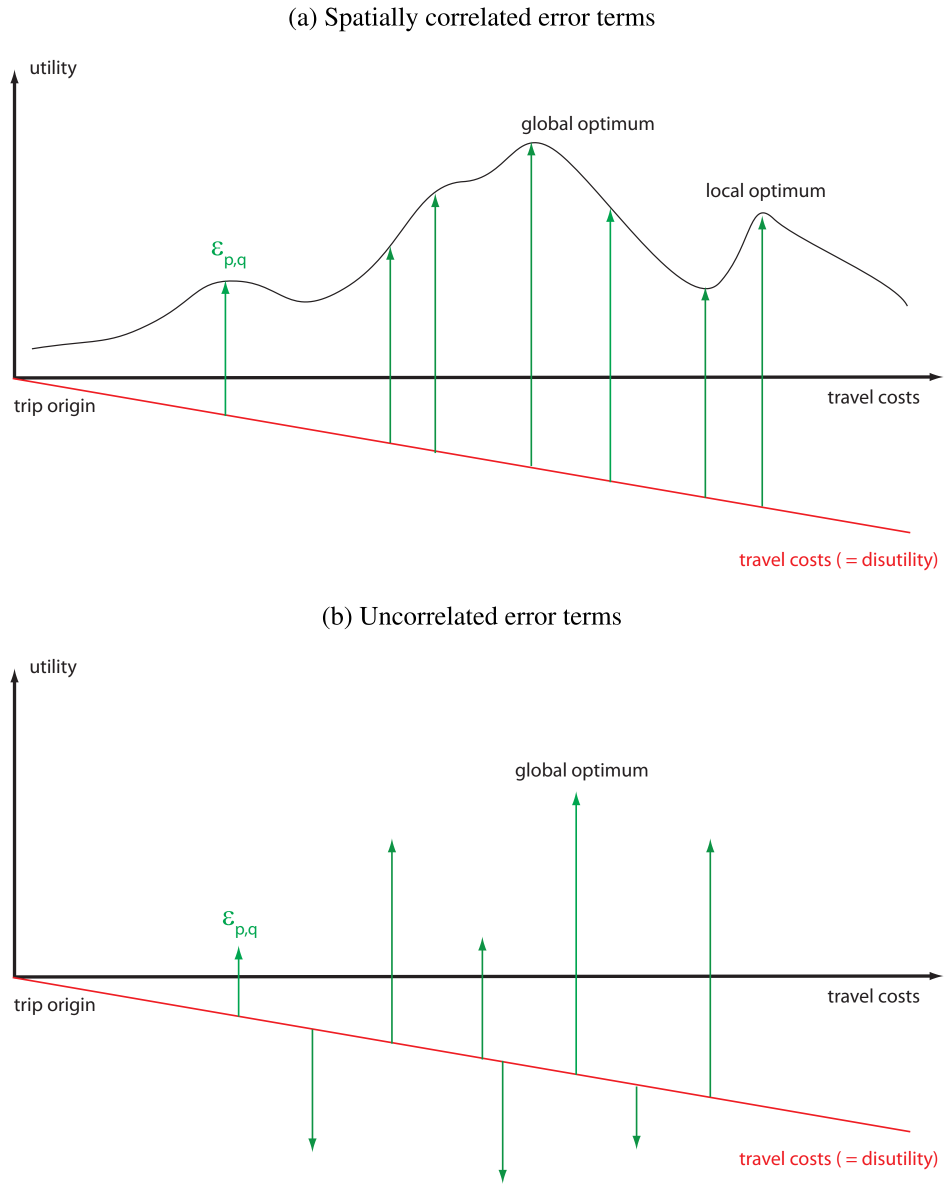

FIGURE 2 Configuration 0 and 1: scores and distances, iteration 200

(a) Configuration 0: scores

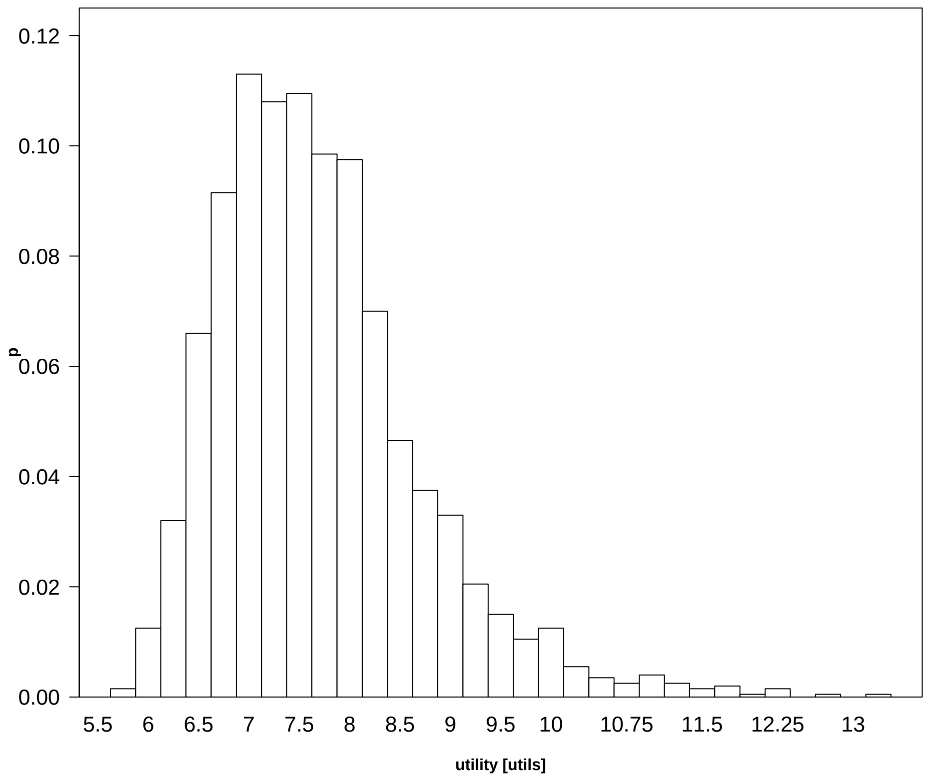

(b) Configurations 0 and 1: shopping distance distribution

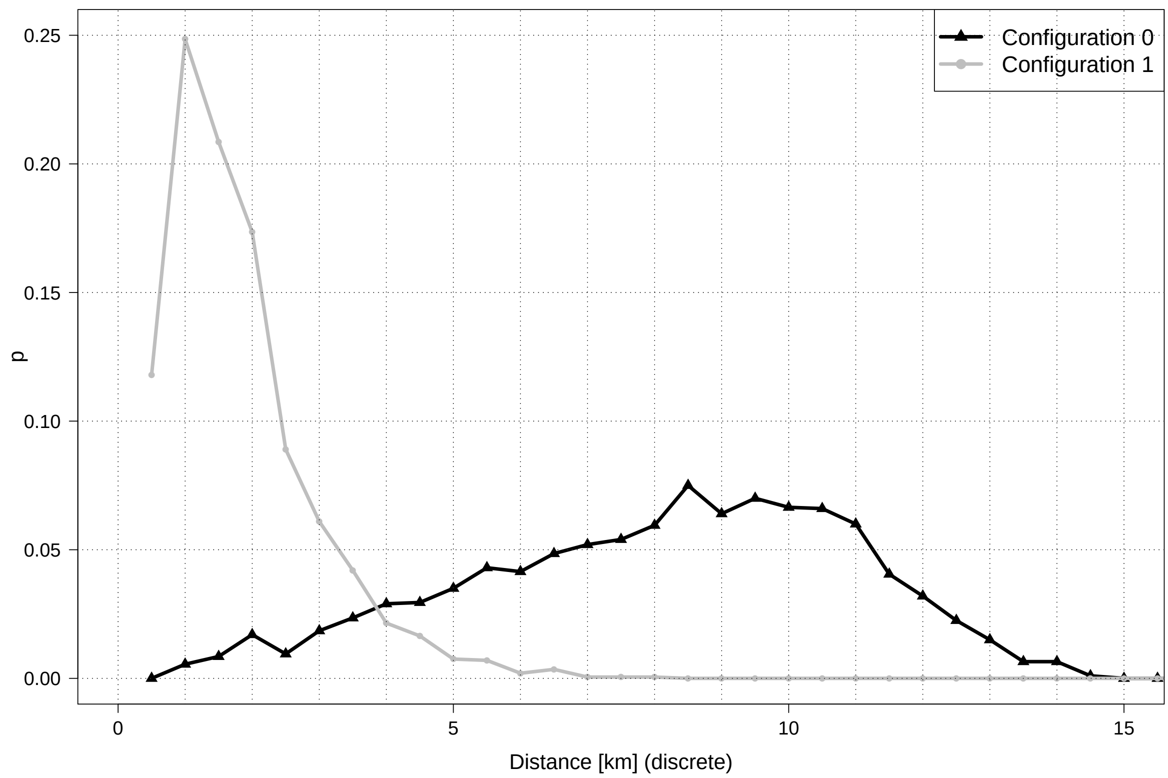

FIGURE 3 Zurich scenario: iteration 100 (relaxed state)

(a) Shopping trips

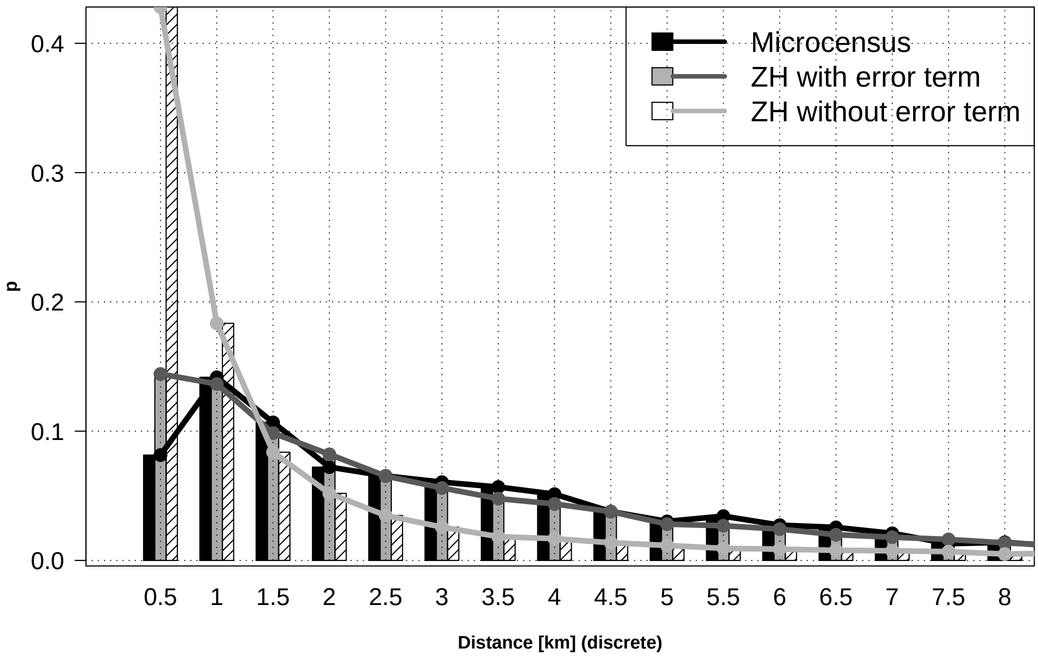

(b) Leisure trips

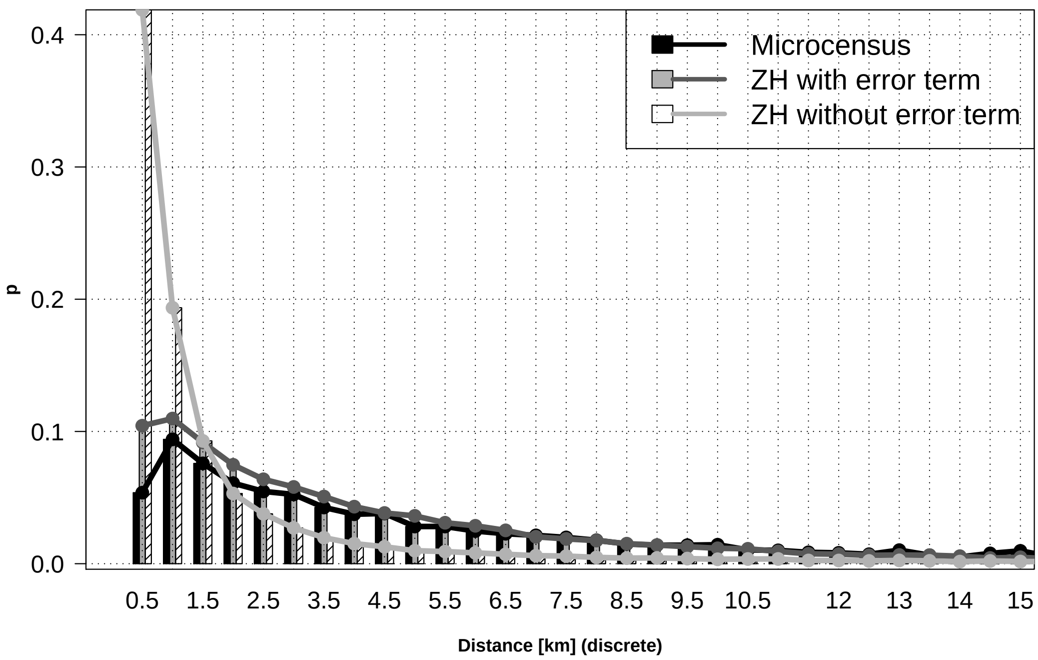

FIGURE 4 Daily traffic volumes for 123 links compared to traffic counts, iteration 100. Per link k the relative error is used, i.e, $100\% \times (vol_{simulated,k} - vol_{counted,k})/vol_{counted,k}$.

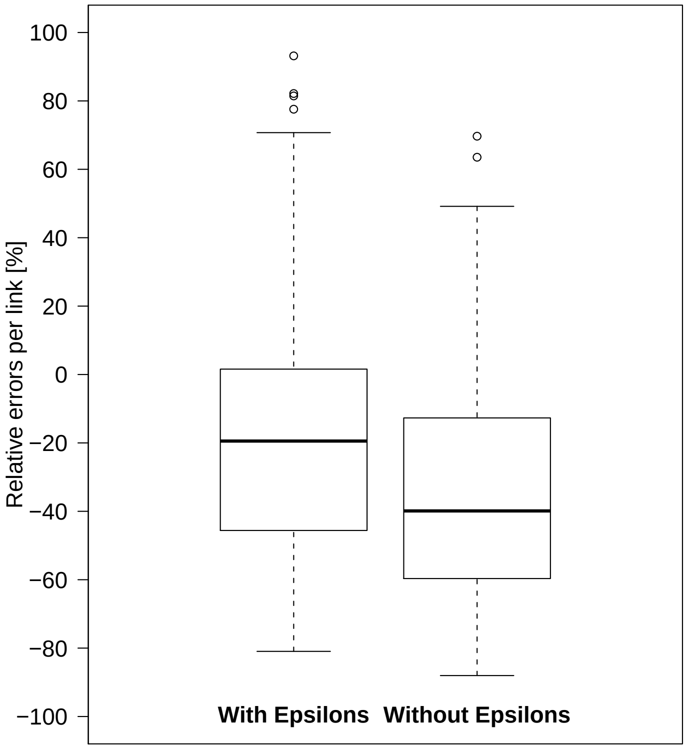

REFERENCES

1. Kitamura, R., A. Kikuchi and S. Fujii (2005) An overview of PCATS/DEBNetS micro-simulation system: Its development, extension, and application to demand forecasting, in R. Kitamura and M. Kuwahara (eds.) *Simulation Approaches in Transportation Analysis: Recent Advances and Challenges*, 371–399, Springer.

2. TRANSIMS (2009) TRansportation ANalysis and SIMulation System, webpage, May 2009, `http://transims-opensource.org/`.

3. Henson, K. M. and K. G. Goulias (2006) Preliminary assessment of activity analysis and modeling for homeland security applications, *Transportation Research Record*, **1942**, 23–30.

4. Timmermans, H. J. P. (2001) Models of activity scheduling behavior, *Stadt Region Land*, **71**, 33–47.

5. MATSim-T (2011) Multi Agent Transportation Simulation Toolkit, webpage, `http://www.matsim.org`.

6. McFadden, D. and K. E. Train (2000) Mixed MNL models for discrete response, *Journal of Applied Econometrics*, **15** (5) 447–470.

7. Hess, S., M. Bierlaire and J. W. Polak (2005) Capturing taste heterogeneity and correlation structures with Mixed GEV models, paper presented at the *84th Annual Meeting of the Transportation Research Board*, Washington, D.C., January 2005.

8. OpenAMOS (2011) OpenAMOS modeling framework, webpage, `http://simtravel.wikispaces.asu.edu/OpenAMOS+Modeling+Framework`.

9. Pendyala, R. M., R. Kitamura, A. Chen and E. I. Pas (1997) An activity-based microsimulation analysis of transport control measures, *Transport Policy*, **4** (3) 183–192.

10. Arentze, T. A. and H. J. P. Timmermans (2000) *ALBATROSS: A Learning-Based Transportation Oriented Simulation*, EIRASS, Eindhoven.

11. Vovsha, P., R. Donnelly and S. Gupta (2008) Network equilibrium with activity-based microsimulation models, *Transportation Research Record*, **2054**, 102–109.

12. Cambridge Systems Inc. (2008) Tel aviv activity schedule travel demand model system: A tour-based approach, *Research Report*, Ministry of Transport, Cambridge Systems Inc., December 2008.

13. Bradley, M. A., J. L. Bowman and B. Griesenbeck (2010) SACSIM: An applied activity-based model system with fine-level spatial and temporal resolution, *Journal of Choice Modelling*, **3** (1) 5–31.

14. Bowman, J. L., M. A. Bradley, Y. Shiftan, T. K. Lawton and M. E. Ben-Akiva (1999) Demonstration of an activity based model system for portland, in H. Meersman, E. van de Voorde and W. Winkelmans (eds.) *World Transport Research*, vol. 3, 171–184, Pergamon.

15. 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, *Transportation Research Record*, **2135**, 87–95.

16. Horni, A., D. M. Scott, M. Balmer and K. W. Axhausen (2009) Location choice modeling for shopping and leisure activities with MATSim: Utility function extension and validation results, presentation, 9th Swiss Transport Research Conference, Ascona, September 2009.

17. McFadden, D. (1978) Modeling the choice of residential location, in A. Karlqvist (ed.) *Spatial Interaction Theory and Residential Location*, 75–96, North-Holland, Amsterdam.

18. Cetin, N. (2005) Large-scale parallel graph-based simulations, Ph.D. Thesis, ETH Zurich, Zurich.

19. Balmer, M., B. Raney and K. Nagel (2005) Adjustments of activity timing and duration in an agent-based traffic flow simulation, in H. J. P. Timmermans (ed.) *Progress in Activity-Based Analysis*, 91–114, Elsevier, Oxford.

20. Lefebvre, N. and M. Balmer (2007) Fast shortest path computation in time-dependent traffic networks, paper presented at the *7th Swiss Transport Research Conference*, Ascona, September 2007.

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

22. Vickrey, W. S. (1969) Congestion theory and transport investment, *The American Economic Review*, **59** (2) 251–260.

23. Balmer, M., A. Horni, K. Meister, F. Ciari, D. Charypar and K. W. Axhausen (2009) Wirkungen der Westumfahrung Zürich: Eine Analyse mit einer Agenten-basierten Mikrosimulation, *Final Report*, Baudirektion Kanton Zurich, IVT, ETH Zurich, Zurich, February 2009.

24. Swiss Federal Statistical Office (2000) Eidgenössische Volkszählung 2000, `http://www.bfs.admin.ch/bfs/portal/de/index/infothek/erhebungen__quellen/blank/blank/vz/uebersicht.html`.

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

26. Swiss Federal Statistical Office (2001) *Eidgenössische Betriebszählung 2001 - Sektoren 2 und 3, GEOSTAT Datenbeschreibung*, Swiss Federal Statistical Office, GEOSTAT, Neuchatel, `http://www.bfs.admin.ch/bfs/portal/de/index/dienstleistungen/geostat/datenbeschreibung/betriebszaehlung05.Document.111424.pdf`.

27. Vrtic, M., P. Fröhlich and K. W. Axhausen (2003) Schweizerische Netzmodelle für Strassen- und Schienenverkehr, in T. Bieger, C. Lässer and R. Maggi (eds.) *Jahrbuch 2002/2003 Schweizerische Verkehrswirtschaft*, 119–140, Schweizerische Verkehrswissenschaftliche Gesellschaft (SVWG), St. Gallen.

28. ASTRA (2006) Swiss Federal Roads Authority, webpage, `http://www.astra.admin.ch/`.

29. Horni, A., D. Charypar and K. W. Axhausen (2011) Empirically approaching destination choice set formation, paper presented at the *90th Annual Meeting of the Transportation Research Board*, Washington, D.C., January 2011.

30. Horni, A., D. Charypar and K. W. Axhausen (2011) Variability in transport microsimulations investigated for MATSim: Preliminary results, paper presented at the *11th Swiss Transport Research Conference*, Ascona, May 2011.