Paper No. 99 09 19

PREPRINT

Microscopic Modelling of Travel Demand: The Home-to-Work Problem

Authors: Peter Wagner and Kai Nagel

Transportation Research Board
78th Annual Meeting
January 10-14, 1999
Washington, D.C.

Los Alamos Unclassified Report LA-UR 98-3368 (This version: December 1, 1998)
Approved for public release; distribution is unlimited

Microscopic Modelling of Travel Demand: The Home-to-Work Problem

Peter Wagner and Kai Nagel

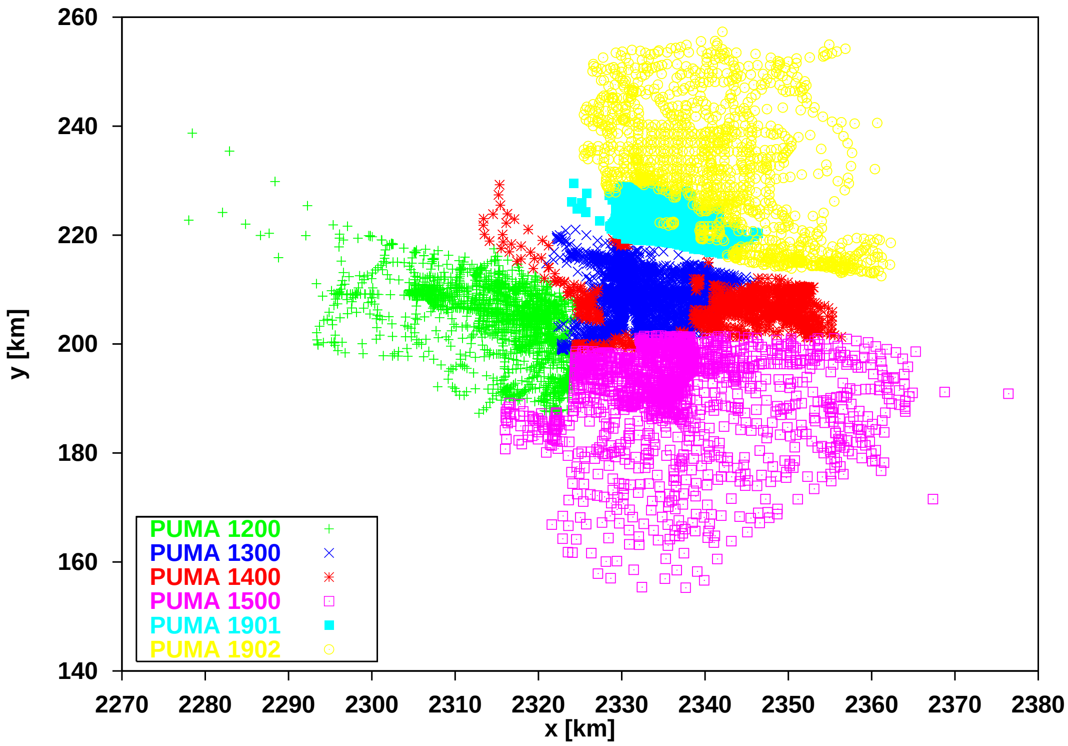

Microscopic Modelling of Travel Demand: The Home-to-Work Problem

Peter Wagner,
German Aerospace Center (DLR)
Porz-Wahnheide, Linder Höhe,
D-51147 Köln, Germany
eMail: peter.wagner@dlr.de

Kai Nagel,
Los Alamos National Laboratory (LANL),
P.O. Box 1663, MS M 997,
Los Alamos, NM 87545
eMail: kai@lanl.gov

This version: December 1, 1998

Abstract

This article describes the results that have been found when tackling the problem of the assignment of employees to their workplaces (destination assignment). The novel feature is the use of a microscopic approach, whose output is suitable for a microscopic traffic simulation. Furthermore, the approach is dynamic since it takes into account departure times (rush-hour) and the jam dynamics in a large transportation system.

This problem is dealt with a microscopic stochastic analogue of the gravity ansatz of transportation planning, described in this article. A microscopic approach has the advantage that there is no need to force the resulting OD–matrix to fulfil the origin and/or destination flows. Any traveler has an origin (home), an associated destination, and returns to home after doing work.

The destination assignment is then used to drive a microsimulation of a transportation system such as TRANSIMS, which consists of an iteration loop to compute the dynamic traffic assignment. The results from the dynamic traffic assignment are then used to eliminate the home-to-work trips that got very long travel times. This creates an additional (outer) iteration loop, consisting of the sequence destination choice – route choice – travel times. This serves (i) to test a very simple destination choice procedure, and (ii) to learn about the interplay between destination choice and route choice in a micro-simulation set-up like TRANSIMS.

Keywords: Activity schedules, assignment of destinations, assignment of routes, micro-simulation of traffic

1 Introduction

A basic object of the traditional four–step–procedure (e.g. [1]) is the origin/destination (OD–) matrix. When using a dynamic microscopic simulation model for computing the time-dependent traffic loads in a network, the OD–matrix has to be decomposed into individual travelers. Therefore, it can be asked if might is more efficient and conceptionally simpler to not start from OD-matrices but from individual trips. A trip is described by an origin, a destination, and a time specification (for example the intended arrival time for the trip). Such trips can be used *directly* in a dynamic traffic assignment, which consists of individual route assignments and a microscopic simulation. Clearly, adding up all such trips over a certain time period leads back to the conventional OD–matrix for that time period.

Yet, generating realistic trips is far from trivial. Most approaches attempt in some way or other to represent aspects of the underlying human decision making process (e.g. [2]). For example, simple gravity models are a functional representation of the plausible assumption that people prefer short trips over long ones, while disregarding many other aspects of the trip decision. Other approaches look at a chain of activities during the day; transportation demand is then generated by the fact that two consecutive activities at different locations need to be connected by a trip. For example, Markov-chain-based models (e.g. [3]) assume that the decision for the next activity is conditioned on the activities performed so far in an activities chain. Models based on utility maximization [4, 5, 6, 7] model human decision-making about their activities as a deliberate maximizing calculation between trade-offs. Some approaches [8] attempt to model humans even more directly, by tracing the actual process that leads to certain schedules. These last approaches can maybe seen as combinations between the Markov-chain-based and the utility maximization approaches. All of these rely on parameters which are functions of socio-demographic variables of the individuals; and the adjustment of these parameters is usually based on one or several surveys.

In this paper, we do not want to join the important discussion which combination of the above models may be the most useful for any given application. Instead, we want to look at the question in how far the new availability of detailed transportation micro-simulation models, which are now technically feasible for large metropolitan areas, possibly modifies the way one can model travel demand generation. In order to make progress on this question, it is useful to start with a subproblem that is as simple as possible, yet still has realistic aspects. The situation that comes to mind here is the travel from home to work, which probably represents a significant amount of morning peak travel. Assuming that the data are available, the maybe most straightforward way to look at this problem is to have a list of workplace locations and a list of worker home locations, and the question is how to match them. This problem is certainly not a new one. For example, "Hitchcock's solution" [1] solves it as an optimization problem, and standard gravity models solve it by making an assumption about the choice probability as a function of distance.

The availability of metropolitan-scale transportation micro-simulations makes it possible to couple the demand generation process to the consequences in the transportation system. As an intuitive example, consistently picking workplaces that are too far from home in the travel demand module will lead to over-saturation of the transportation system, which will make it impossible for *everybody* to execute their plans. In itself, this is probably not a very surprising approach to practitioners in the field – it seems to be common practice in the four-step-process to adjust the OD-matrix to make traffic in the assignment more realistic. Given our new simulation capabilities, we would like to, however, eventually reach a point where we can make such adjustments both mechanical and based on individual travellers. "Mechanical" because we would like a computer to do the adjustments; and based on individual travellers because we believe that longterm progress can be made by having individual travellers accessible as individual objects in the methodology.

In this paper, after describing our data (Sec. 2), we will describe an approach that makes very simple assumptions about the workplace assignment procedure (Sec. 3.1). We then describe how the results of this assignment are fed into a dynamic traffic assignment which consists of several iterations between a routing module and a traffic micro-simulation (Sec. 3.2). We then describe how the result of the dynamic traffic assignment can be used to adjust the workplace assignment (Sec. 3.3). Last, we compare our results from this extremely simplified approach against some Portland data (Sec. 4). The paper is concluded by a discussion and a summary.

Note that in this paper, no attempt is made to systematically match ("calibrate") simulation results against the Portland data. Such a study is, however, the subject of current research.

2 Study context

The work described here depends on the availability of an excellent data-basis, being the case for Portland, which is the current focus of the TRANSIMS project [9]. However, one can hope that similar data will be available for other cities in the near future as well.

2.1 Input data

This study has used the following input data:

- A reduced road network of Portland with 8,564 nodes and 20,024 links. This road network is actually a version of the network that Portland has traditionally used for their four-step-process (Bill Stein, personal communication).

- Data from the census in 1990 which were input to an algorithm [10] to generate synthetic households with 1,415,900 household members, of which roughly 660,000 are employees.

- A detailed data-set containing the locations of the employers of Portland with the approximate number of workplaces.

The workplaces as well as the employees (given by the synthetic households) are associated with the links in the network.

2.2 Comparison data

The Portland metropolitan area stretches across two states, Oregon and Washington. For the part in Washington State (Vancouver), employer information was not available. The missing information was generated by counting the numbers of all necessary workplaces from the

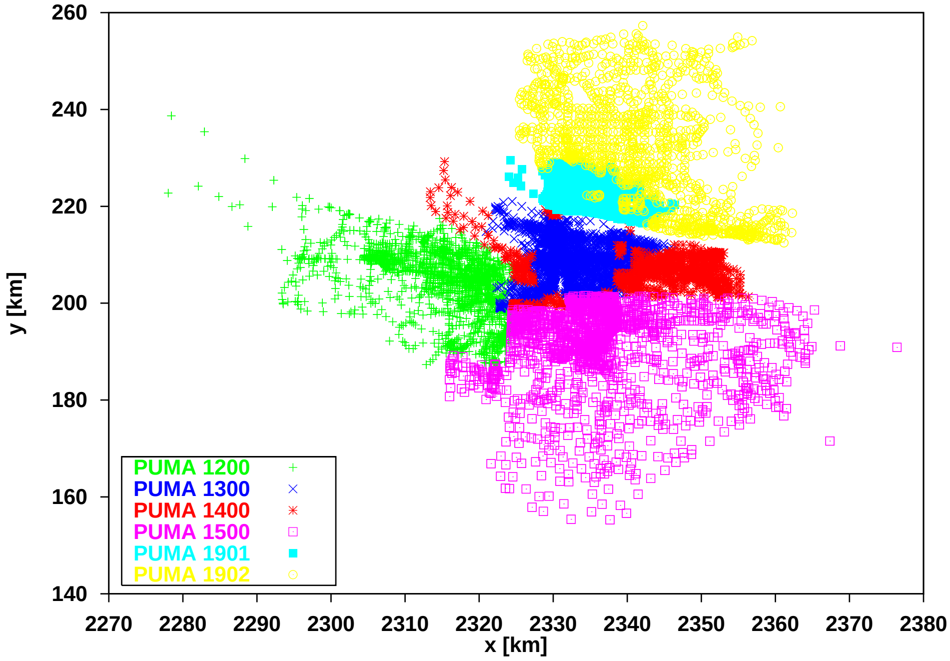

Figure 1: Spatial decomposition of Portland into PUMA's. The numbers in the legend are the PUMA-ids.

demographic data (as given by the synthetic household data) and subtracting the number of workplaces in Oregon. The remaining workplaces were randomly assigned to links in the area of Washington State in proportion to the population density. The synthetic household data additionally contain the information in which PUMA (Public-Use Microcensus Area, see Figure 1) the employees are working. Although the spatial decomposition of Portland into the PUMA's is a very rough one, this information will be used in order to judge the quality of our results.

The synthetic household data contain additional information about when people start their work, and their actual trip duration. The appropriate description of this duration is a frequency distribution, shown in Figure 2. There the number of persons $p(\tau)$ that need a certain time τ to go to work is plotted. As usual, such a distribution from a survey is clustered around values as "5 min", "10 min", "15 min", "30 min", etc. In order to utilize this distribution in a computational method, it needs to be smoothed to second-by-second values. For this purpose, we did a fit by a beta-distribution:

$$p(\tau) = \mathcal{N}\, \tau^{\alpha} (\mathcal{T}_{\max} - \tau)^{\beta}, \tag{1}$$

where $\mathcal{N}$ is a normalization constant. The parameters of this distribution cannot determined directly from $p(\tau)$, they have been obtained by a non-linear least-squares fit of the integrated distribution to the integral of a beta-distribution $P(\tau)$:

$$P(\tau) = \int_0^{\tau} p(\tau')\, d\tau' \propto \int_0^{\tau} \tau'^{\alpha} (T_{\max} - \tau')^{\beta}\, d\tau'. \tag{2}$$

The parameters from the fit are $a = 0.83$, $b = 6.03$, and $T_{\max} = 100[min]$.

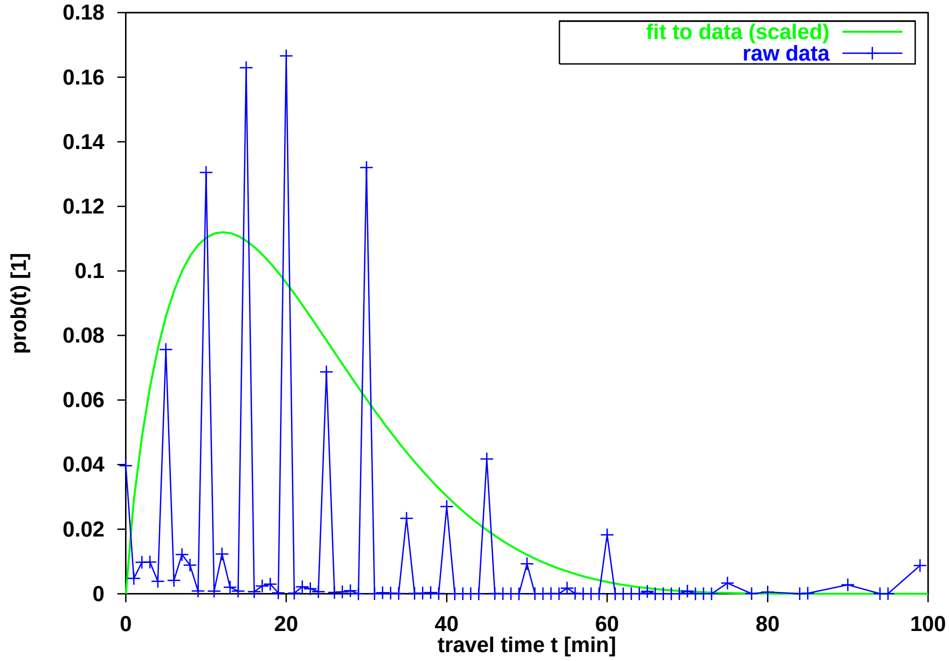

Figure 2: Empirically determined travel time distribution, together with a smooth approximation obtained by the procedure described in the text; note the peaks at multiple integer of 5 [min], due to the tendency of humans to report simple numbers only.

Data such as this has to be used with care, though. There is widespread suspicion that the travel times reported in the census are longer than the one people are actually on the way. One possible explanation for this is that people report the time they "plan" for their trip, not the time they are actually in the transportation system (K. Lawton, personal communication, and a referee).

3 Our approach

3.1 Assignment of workplaces

As already described, the data-set available for Portland contains information about the number of workplaces W_i at each spatial index i. The spatial indices are associated with the links of the network. The synthetic household data contain the number of employees E_i, where again the index i refers to a spatial location. What has to be computed is, for any of the employees given, her or his respective workplace. We can for example make the standard assumption that, for each employee n, there is a choice function, $P_n^{choice}(\tau)$, which will give us that probability that this employee will accept a travel time of τ. Such an assumption can for example be derived from utility theory (e.g. [4]), although in this paper, we do not make any a priority about its functional form. For the sake of methodological progress, we will, however, assume that the choice function is the same for everybody, i.e. $P_n^{choice}(\tau) \equiv q(\tau)\ \forall n$. This is clearly a strong assumption and in violation of empirical findings (e.g. [5]), but we want to accept it in this first attempt to include microsimulation feedback into the demand generation process.

As mentioned above, a number of households as well as a number of workplaces is associated with each link in the Portland network. Labeling these links by an arbitrary index i, the assignment proceeds by performing the following steps for any i:

1. Attractions $a_{j,j \neq i}$ are computed for any of the other links with $a_j = W_j\, q(t_{ij})$ where W_j is the number of workplaces at link j, t_{ij} are the travel times based on the shortest paths in the network and $q(\tau)$ is the choice function.

2. The a_j are interpreted as probabilities that this link j is chosen, so any of the employees E_i in link i is associated according to these probabilities to one of the workplaces.

3. After assigning a person from i with a workplace in j, the number of workplaces still available for assignment is reduced by one.

4. After having assigned any of the employees in i, the next link to process is drawn at random and the algorithm continues with step 1.

Since the number of workplace left for assignment decreases during this process, in the end all employees are assigned to all workplaces. This book-keeping ensures that the constraints on the OD–matrix $F_{i \to j}$ are fulfilled:

$$E_i = \sum_j F_{i \to j}, \tag{3}$$

$$W_j = \sum_i F_{i \to j}. \tag{4}$$

3.2 Route assignment and traffic micro-simulation

Trips generated by the above home-to-work trip generation module are then fed into a router, which assigns fastest paths to all of the trips. These fastest paths are computed by using the simulated time-dependent link delay times; for the initial routing, when simulated link delays are not yet available, free speed link delay times are used. For details on the Dijkstra algorithm used, see for instance [11, 12]. The routes are input into a micro-simulation that executes these routes in the network, giving "real" travel-times for all the trips. These real travel times are then used to re-route a certain part of the travelers, which gives a new set of routes which are executed, until this process has converged.

The micro-simulation used for this paper is a queue-type simulation [13, 14] which describes the main effects relevant for the determination of travel times in a dynamically driven network: (i) the spill-back phenomenon of a loaded link, (ii) the FIFO-condition (First-In, First-Out, meaning that a car which enters a link first leaves this link first), and (iii) that the travel times that depend on the load of a link. Any link in the network is modeled as a waiting queue with an exit flow constraint and a storage constraint. Cars are allowed to enter the link only if there is enough free space; otherwise, they wait on the link where they are coming from. This kind of simulation gives an approximation to the traffic load distribution in the network and has the advantage that it is computationally very fast. Running the queue-model for the morning period, lasting from 4 a.m. to 1 p.m. with 520,000 cars, needs 25 min on a SUN Ultra 250 MHz.

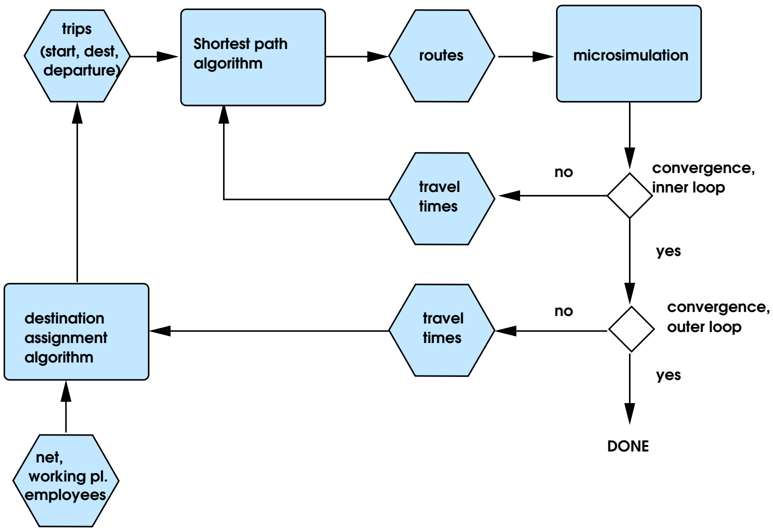

Figure 3: Flow diagram of the simulations used to compute the combined destination and route choice assignment. The inner loop is the route assignment loop, the outer loop computes the destinations based on the results of the route assignment.

3.3 Iterative adjustment of workplace locations

Whichever choice function $q(\tau)$ we select, we want our procedure to be consistent, in the sense that the travel times assumed in the trip generation module should be consistent with the travel times that the simulation generates. This can be solved by using a second iterative loop outside the route assignment loop described in the last section (see Figure 3). In order to distinguish the assignment of destinations from the route assignment of the dynamic traffic assignment algorithm, the term "destination assignment" will be used in the following. The whole algorithm then looks as follows:

(1) Assign initial workplace destinations. Travel time from a home location i to a potential work location j is computed using a fastest path algorithm assuming free speeds on the links.

(2) Perform a dynamic route assignment for these trips as described in Sec. 3.2. This means as many routing steps and micro-simulation runs as necessary until relaxation is reached.

(3) Re-assign workplace destinations for part or all of the population assuming the link delay times generated by the micro-simulation.

(4) Goto (2) except when results do not change any more.

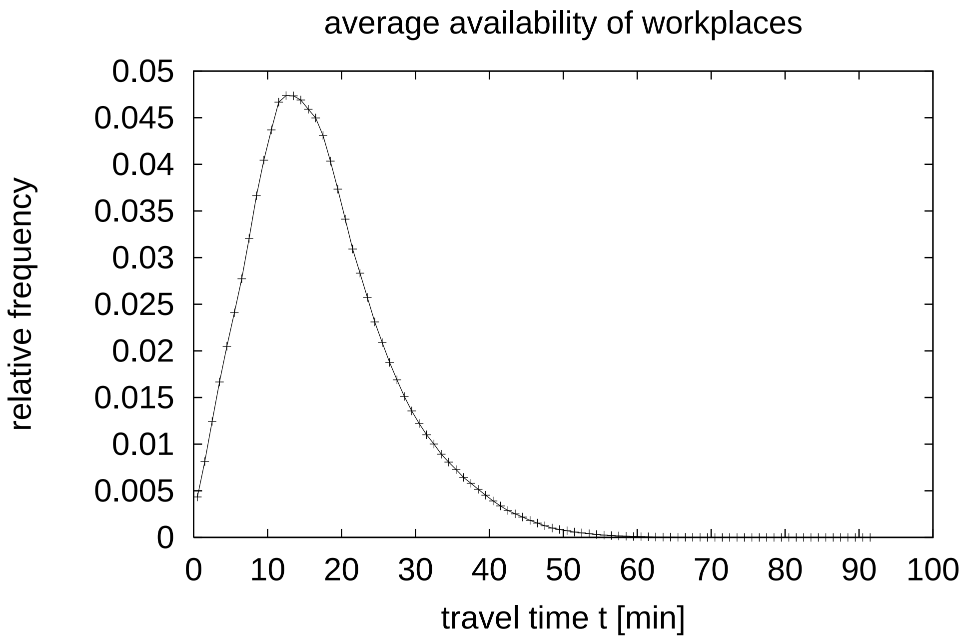

Figure 4: Average workplace availability. Probability, averaged over all worker home locations, that there is a workplace at travel time t. Empty network.

4 Simulation Results

4.1 Specification of the choice functions and initial workplace assignment based on empty network

For any possible home location i one can write down a function

$$N_i(T) = \sum_j W_j \, \delta(t_{ij} - \tau) \ ,$$

which denotes the number of workplaces available at travel time τ. Remember that W_j is the number of workplaces on link j and t_{ij} is the travel time from link i to link j. $\delta(t_{ij} - \tau)$ is an indicator to take the sum only over the workplaces at the "correct" time distance.

Summing this over all employers, one obtains an "average" workplace availability function:

$$N(\tau) := \frac{\sum_i E_i \, N_i(\tau)}{\sum_i E_i} \ .$$

This function is plotted in Fig. 4, where the travel times T are computed using fastest path on the empty network.

Now, when every worker would randomly select his/her workplace, and the resulting travel could be done at free speeds, then the resulting travel time distribution $p(\tau)$ would just look like $N(\tau)$. This makes it tempting to just try $q(\tau) = const$ as a choice function, i.e. to select a workplace at random, independant of the travel time. Fig. 5 (curve labelled "flat attraction")

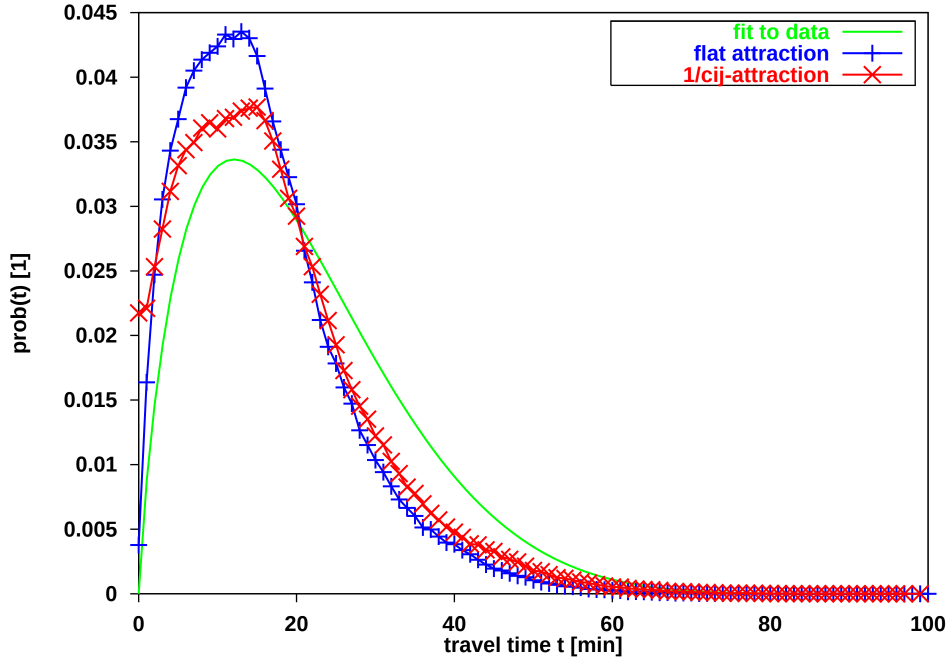

Figure 5: Resulting travel time function from different assignment cost functions, compared to the empirical function (fit to data). The $1/c_{ij}$–attraction function results from an destination assignment with the t^{-1}–function, flat attraction is the function without distance dependence mentioned in the text. These are travel times expected by the initial destination assignment without having executed the micro-simulation.

shows the result of such an assignment for the empty network. Deviations from the curve for $N(\tau)$ can be explained by the "bookkeeping" piece of the code (see text before Eq. 3). Such an assignment can for example be used as a benchmark; note that discrete choice models often note the likelihood of getting the survey from totally random choices as a benchmark.

For comparison, Fig. 5 also shows how the same assignment would look if $q(\tau) = 1/\tau$ would be used as choice function (curve labelled "1/cij attraction"). The unexpected fact that this choice actually generates *more* trips with durations of 20 min or more is probably due to the fact that the close destinations get "used up" by the first employees to be assigned, and so later during the assignment, only destinations far away are available. This choice function is shown for comparison purposes and will not be used in the remainder of this paper.

4.2 Iterated results

After 25 route iterations, the simulated $p(\tau)$-distribution did not change any more, and thus the route assignment iterations based on the initial destination assignment were stopped. Travel times for the initial empty network and the final loaded network were computed by a Dijkstra algorithm for time-dependent link-costs. These two travel times, available for each traveler, were used for re-assigning a number of employees. Only the travelers whose travel-time in the loaded network exceed the travel-time in the empty network by a certain amount θ were subject

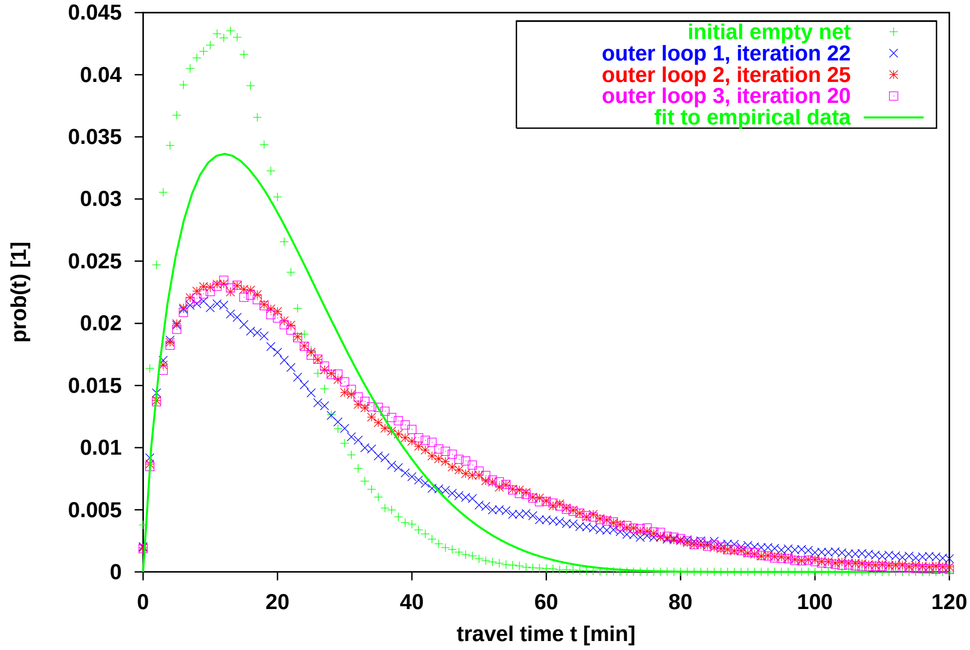

Figure 6: This plot shows the evolution of the travel time distribution during the assignment loops. Shown here is the fit to the empirical travel time distribution, the initial assignment based on the empty net, and three intermediate results during the iterations.

to re-assignment. The difference between the two travel times was measured by computing

$$\eta = \frac{T_{\text{full}} - T_{\text{empty}}}{T_{\text{full}} + T_{\text{empty}}}$$

which maps the difference in travel times to the interval $[-1, 1]$. The re-assignment algorithm is very simple: the workplace of any traveler who matches the condition $\eta > \theta$ above was given back to a pool of workplaces, and then the procedure described above that assigns the employees to the workplaces was repeated for this subset of all travelers. This yields a new set which was then fed into the iteration that determines the dynamic equilibrium. The whole process has to be repeated until all travelers are satisfied with their choice, or until there are

O/D–pair	Census data	Assignment
O $\to$ O	423,362	393,942
O $\to$ W	8,282	37,702
W $\to$ O	29,732	51,864
W $\to$ W	48,971	26,839
Total trips	510,347	510,347

Table 1: Comparison between simulation and data on the basis of an aggregated OD-matrix. The flows between Vancouver (W) and Portland (O) are shown.

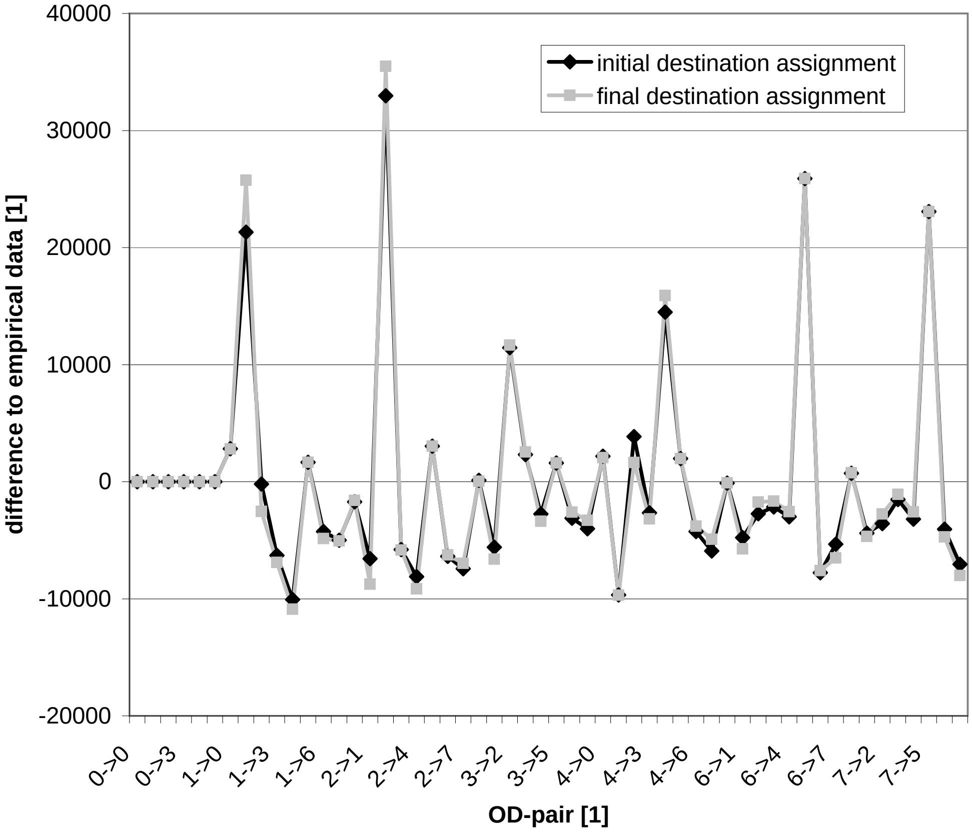

Figure 7: Comparison of the results of the simulation with the empirical data. The left plot shows a comparison between empirical data and the simulation for the OD-streams between the Portland PUMA's. In this case the re-assignment of the destinations didn't improve the original assignment.

no more changes between subsequent iteration steps.

What this means is that the model assumes that travelers do not care how long they drive to work as long as their travel time is not too much longer in the congested network than it would be in the empty network. It is however, not a priori clear if it is possible for the algorithm to re-arrange all trips in a way that all travelers are "satisfied" according to this criterion. The stopping criterion for the destination assignment iteration will thus be that, for a given θ, the *same* fraction of travellers is unsatisfied than in the last destination assignment iteration.

The first re-assignment of the destinations was done with $\theta = 0.5$, leading to 150,000 out of 520,000 new destinations. Then the route assignment was restarted, again with an initially empty network and the corresponding travel times. After the second route assignment, with $\theta = 0.4$, roughly 100,000 travelers (i.e. 20%) got new destinations. After the third outer iteration loop was completed, again roughly 20% of all travelers have travel times that are larger than 40% of their expected travel times, so the iteration has met the above defined termination criterion.

One should probably note that the way home locations and work locations are matched is not behaviorally realistic. In practice, people probably adjust for example shopping locations according to where they live and where they work. However, the choice of the home location may follow the choice of the work location more often than the other way around. The reason why we do it this way around in our current work is because all the demographic information we have about the worker is coupled to the home location. We cannot just move workers with given workplaces to different home locations without taking precautions that we do not change the demographic characteristics of our region away from the data that we have.

The results of the simulation are summarized as follows. In Figure 6, the measured travel times are compared to the simulated ones. One sees that the initial assignment assumption, based on the empty network, assumes too many trips with durations of about 18 minutes, and not enough around 40 minutes. After executing the simulations, this tendency is reversed. Also, after the first iteration, the algorithm actually reduces the number of very long trips and increases the number of trips around 30 minutes. One could speculate if a "tighter" choice of θ could increase this trend. Overall, one can maybe say that the benchmark choice of a "flat" choice function results, maybe unexpectedly, in a travel time distribution that is not completely unrealistic. Nevertheless, more work needs to be invested to make these results more realistc.

Table 1 shows the traffic flows between Vancouver (Washington, W) and Portland, Oregon (O)), compared to the data obtained from the census data. The assignment has an average error of about 15 % . No attempt to fit the simulation to the data has been done. The only parameter that determines the results here is the re-assignment factor θ. The simulation results basically assign too much traffic between O and W, and they only partially catch the asymmetry between O and W. Different realizations (different random number seeds) of this stochastic assignment algorithm lead to similar results, where the statistical errors are of the order of what has to be expected naively (i.e. any number in this table has an error of the order of the square-root of this number). Therefore, the algorithm is robust against small changes in the parameters. In Figure 7, a more detailed comparison between the measured OD–matrices (between the Portland PUMA's) is shown. In order to facilitate the comparison, the difference between the empirical OD–matrix and the simulated ones is shown, both for the initial destination assignment as well as for the third reassignment. It seems, however, that the reassignment did not modify the structure of these OD flows very much.

Note finally that the algorithm that computes the destinations runs very fast even when using the exact distances in the network as computed by the Dijkstra algorithm. The assignment of 600,000 individuals took about 20 minutes of computing time. When using the flat travel time function which does not need Dijkstra, the main time is spent with reading and writing the results from/to disk.

5 Summary

It has been shown that it is possible to use micro-simulation techniques for a feedback process regarding travel demand generation. Specifically, we have shown that an outer iterative loop can be performed outside the typical dynamic route assignment iterations. As a benchmark test, we have used a random selection of workplaces given home locations of workers. Maybe surprisingly, this simple choice leads to traffic and travel time distributions that are not as unrealistic as one might expect. The perhaps most unsatisfactory aspect of this approach is that it does not result in a clear criterion according to which one should do the workplace location re-assignment because the workplace selection does not depend on travel time in the first time. Further work into this direction is currently being performed at TRANSIMS, both in order to enhance this simplified approach to be somewhat more realistic, and in order to include sophisticated activities generation modules into TRANSIMS.

Acknowledgements

It is a pleasure for one of us (PW) to thank Los Alamos National Laboratory (LANL) for the kind hospitality during a stay in Los Alamos, where part of this work was conducted. Also thanks to colleagues in the research cooperative for simulation of traffic and impacts on the environment (FVU) for numerous discussions about the topic presented here. Special thanks to Guido Rindsfüser from the RWTH Aachen, who has helped a lot in understanding the classical approach on this theme.

References

[1] Y. Sheffi. *Urban transportation networks: Equilibrium analysis with mathematical programming methods.* Prentice-Hall, Englewood Cliffs, NJ, USA, 1985.

[2] D. Ettema and H. Timmermans, editors. *Activity-based approaches to travel analysis.* Pergamon, 1997.

[3] R. Kitamura, S. Fujii, A. Kikuchi, and T. Yamamoto. An application of a micro-simulatur of daily travel and dynamic network flow to evaluate the effectiveness of selected tdm measures for CO_2 emissions reduction, 1997.

[4] M. Ben-Akiva and S.R. Lerman. *Discrete choice analysis.* The MIT Press, Cambridge, MA, 1985.

[5] John L. Bowman. *The day activity schedule approach to travel demand analysis.* PhD thesis, Massachusetts Institute of Technology, Boston, MA, 1998.

[6] A system of activity-based models for Portland, Oregon, Draft final report, 1997.

[7] W.W. Recker. The household activity pattern problem: General formulation and solution. *Transportation Research B*, 29:61–77, 1995.

[8] T.A. Arentze, F. Hofmann, C.H. Joh, and H.J.P. Timmermans. Experiences with deveoloping ALBATROSS: A learning-based transportation oriented simulation system. In *Verkehr und Mobilität*, number 66 in Stadt Region Land, pages 61–70. Institut für Stadtbauwesen, Aachen, Germany, 1998.

[9] TRANSIMS, TRansportation ANalysis and SIMulation System, Los Alamos National Laboratory, Los Alamos, U.S.A. See `transims.tsasa.lanl.gov`.

[10] R.J. Beckman, K.A. Baggerly, and M.D. McKay. Creating synthetic base-line populations. *Transportion Research Part A – Policy and Practice*, 30(6):415–429, 1996.

[11] C.L. Barrett, R. Jacob, and M.V. Marathe. Formal language constrained path problems. Los Alamos Unclassified Report (LA-UR) 98-1739, Los Alamos National Laboratory, Los Alamos, NM, U.S.A., http://transims.tsasa.lanl.gov, 1997. to be published SWAT.

[12] R.R. Jacob, K. Nagel, and M.V. Marathe. A computational study of routing algorithms for realistic transportation networks. Los Alamos Unclassified Report (LA-UR) 98-2249, Los Alamos National Laboratory, http://transims.tsasa.lanl.gov/, 1998.

[13] C. Gawron. An iterative algorithm to determine the dynamic user equilibrium in a traffic simulation model. *International Journal of Modern Physics C*, 9(3):393–407, 1998.

[14] Patrice Simon and Kai Nagel. Simple queueing model applied to the city of Portland. Los Alamos Unclassified Report (LA-UR) 98-9303, Los Alamos National Laboratory, http://transims.tsasa.lanl.gov/, 1998.

[15] Los Alamos Unclassified Report (LA-UR), Los Alamos National Laboratory, http://transims.tsasa.lanl.gov/.