Los Alamos Unclassified Report LA-UR 98-4389
(This version: January 24, 2000)
Approved for public release; distribution is unlimited.

TRANSIMS for urban planning

Kai Nagel, Richard J. Beckman, and Christopher L. Barrett

TRANSIMS for transportation planning

Kai Nagel, Richard J. Beckman, and Christopher L. Barrett

Los Alamos National Laboratory, Mail Stop M997,
Los Alamos NM 87545, U.S.A.

kai@santafe.edu, barrett@tsasa.lanl.gov, rjb@lanl.gov

January 24, 2000

Abstract

The TRANSIMS (TRansportation ANalysis and SIMulation System) project at the Los Alamos National Laboratory attempts to model all aspects of human behavior related to transportation in one consistent simulation framework. Input to TRANSIMS are transportation infrastructure data, demographic data, land-use data, and knowledge about human decision-making. Output data are any set of traditional or non-traditional measures of effectiveness (MOEs) that one could obtain from a second-by-second knowledge of the transportation system.

The key to TRANSIMS is a completely microscopic simulation of the travelers. Demographic data is disaggregated into synthetic individuals; these individuals make plans about their activities (work, sleep, eat, etc.) and how they intend to get their; finally, a realistic transportation microsimulation executes all plans simultaneously and thus computes the interactions between plans, for example resulting in congestion. Iterating between plans-making and plans-execution (the microsimulation) simulates day-by-day learning.

This text gives an overview of the current status of TRANSIMS, as well as pointers to detailed studies that have been done using it.

1 Introduction

In the context of urban planning it would be useful to have a computational tool that evaluates transportation consequences of urban evolution scenarios. The anticipated urban structure, including anticipated demographic data and anticipated transportation infrastructure, could then be fed into a computer, and the computer would calculate the resulting traffic. Maybe the tool would even have a virtual reality component, enabling one to zoom to one's home and count the number of passing cars during the rush hour (e.g. Fig. 1). Clearly, such a tool should include the effects of "induced" traffic, i.e. the observation that lower congestion levels encourage people to more travel; the tool should be able to say something about the variability of traffic; and it should be capable of incorporating new technology, such as telecommuting or telematics systems.

The problem of induced traffic makes clear that transportation is *not* a simple infrastructure problem, where to find a good or optimal solution to move a given demand, but a problem where individual people's values and preferences play an enormously important role. In consequence, any quantitative technique needs to be able to represent aspects of human decision-making. The problem thus becomes as much a problem of (computational) social science as one of engineering.

The TRANSIMS (TRansportation ANalysis and SIMulation System) project at Los Alamos National Laboratory (TRANSIMS,1999) is an attempt to build such a tool. The key to the TRANSIMS design is that it is completely microscopic, which means that it keeps track of individual travelers throughout its modules. Similarly, elements of the transportation infrastructure, such as intersections, traffic lights, turn pockets, etc., are represented microscopically.

In TRANSIMS, each traveler is a computational agent. Agents make plans about what to do during a day – in order to get from one activity location to another, agents can, for example, walk, use bicycles, drive cars, or use busses. Eventually, all plans are simultaneously executed in a micro-simulation of the transportation system.

In principle, this leads to a straightforward simulation approach (see Fig. 2): Derive synthetic households from demographic data and locate them on the network; use the demographic information together with land use information to derive activities (working, sleeping, eating, shopping, etc.) and activities locations for each household member; and let agents decide about mode and routing for their transportation. So far, all these are *plans*, i.e. *intentions* of the simulated individuals.

These plans can then all be fed into a realistic transportation micro-simulation, which can be used as the basis for analysis, such as emissions calculations.

The advantage of such a microscopic approach is that, at least conceptually, it can be made arbitrarily realistic. This makes it possible to include dynamic effects such as queue spillover, which are sometimes hard to represent in traditional methods. It also makes it possible to include new and perhaps unanticipated technology at a later time. For example, the whole architecture of ITS (Intelligent Transportation Systems) can be mirrored by a careful software implementation.

Yet, there are also several disadvantages, some of them being:

(i) Size of the problem: Metropolitan regions typically consist of several millions of travelers. Executing a second-by-second transportation micro-simulation on a problem this size within reasonable computing time is only possible with the use of advanced statistical and computational techniques.

(ii) Behavioral foundation: We are far from understanding human behavior. For that reason alone, we are unable to predict the behavior of *individual* travelers. However, there seems to be a realistic chance that the *macroscopic* (emergent) behavior that is generated by thousands or ten-thousands of interacting individuals is considerably more robust than the behavior of an individual agent. This would be similar to Statistical Physics, where the trajectory of a single particle is unpredictable, yet, useful macroscopic properties of gases such as equations of state can still be derived.

(iii) Consistency problem: The approach outlined above is not as straightforward as it sounds because the plans depend on *expectations* about traffic conditions during execution. For example, if a person expects congestion, he or she may make different plans than when no congestion is expected. Yet, congestion occurs only when plans *interact* during their simultaneous execution. In short, plans depend on congestion, but congestion depends on plans. – This logical deadlock is not unknown in economic theory and is traditionally overcome by the assumption of rational agents. Both with and without the assumption of rationality, this problem of consistency between plans and micro-simulation makes the computational challenge even bigger.

(iv) Robustness: Any approach to a problem needs to have reproducibility of the results under a wide enough range of changes, or otherwise the results are useless for practical purposes.

In the remainder of this paper, we want to give information on how the basics of such a transportation simulation can look like. Because of our own experiences, we will focus on the TRANSIMS project; however, there are other projects in this area (e.g. FVU-NRW, 1999).

2 TRANSIMS Modules

Disaggregation of demographic data – The idea behind the TRANSIMS approach is that it simulates *individual* travelers. However, the typically available demographic input data is aggregated over geographical zones. The synthetic population generation module (Beckman et al, 1996) of TRANSIMS generates instances of stochastic populations. The synthetic population can include information about household membership, home location, income, vehicle availability, etc.

Generation of individual plans – In TRANSIMS, agents make plans about their daily behavior. Individual characteristics of the agents (e.g. income, car ownership) are used together with land use information to derive activities (working, sleeping, eating, shopping, etc.) and activity locations for each individual agent. This can for example be done via complicated activities models with hundreds of free parameters obtained from the estimation of surveys (e.g. Bowman, 1998), or via simplified "toy" models for research purposes (e.g. Esser and Nagel, 1999). Such models should include activity *location* because that is where the travel demand is derived from.

Activities at different locations need to be connected by transportation. Synthetic travelers need to select their modes (walk, drive, transit, etc.), and their routes. Again, the methods can range from simple, Dijkstra-type fastest-path algorithms (Jacob et al, in press) to sophisticated formal language constrained path problems (Barrett et al, 1997).

The result of the plans generation module is a file, which contains, for each agent in the simulation, the blueprint for the day, including all planned activities and their locations, and a detailed description (including information such as the route or the bus lines) of how she/he plans to get from one activity to the next.

Micro-simulation – All agents' plans are executed simultaneously in a transportation micro-simulation. This is the only place where interactions between agents are actually computed. Micro-simulations, too, can reach from relatively simple queue models (Simon and Nagel, 1999) to arbitrarily realistic implementations (e.g. Krauss, 1997; MITSIM, 1999); the currently maybe most realistic implementation that still fulfills the computing requirements is based on a cellular automata technique for car following and lane changing, enhanced by additional rules for elements such as signals, unprotected turns, turn pockets and weaving lanes, etc. (Nagel et al, 1997).

Analysis – Such scenarios can be analyzed in detail in any way one wishes. Note that, because of the nature of the microscopic approach, all microscopic variables are still accessible at this stage. For example, information can be separated out by income of the travelers, by trip purpose, by age, etc. Since one can run the simulations for different alternative scenarios regarding the transportation infrastructure, it is straightforward to for example consider the effects of certain infrastructure alternative on certain sub-populations (Beckman et al, 1997; see Fig. 3). Also, since driving is –at least in principle– realistically modeled, vehicular emissions can be calculated (Williams et al, 1998).

3 Feedback and re-planning

As stated in the introduction, for recurrent situations such as rush-hour traffic, the conditions *expected* during plans-making need to be consistent with conditions *encountered* during the execution of the plans. If they are not consistent, people are likely to change their plans. For example, if congestion turns out to be much worse than expected when the plans were made, people may drop less important destinations from their trip. That means that a situation where expected conditions are not consistent with conditions during execution are transient. Since we are currently interested in "average" (i.e. non-transient) results, such inconsistency is undesirable.

Yet, when the plans are made for the first time, no information about the resulting dynamics is available. This is true both for models and for reality. In our simulations, we currently solve this problem by a relaxation method, i.e. by iterating between the plans-making modules and the micro-simulation until expectations during plans making and conditions during plans execution are consistent (Nagel and Barrett, 1997; Rickert, 1998; Rickert and Nagel, 1998; Esser and Nagel, 1999; see Fig. 4). ,

One should note here that there are many questions that need to be answered for such an approach. For example: What are *good* indicators of relaxation? Do different initial conditions and/or iteration schemes lead to different relaxed states ("path dependence")? If not, can one speed up the relaxation process? Is reality at all similar to the relaxed state obtained with this methodology? If not, is there a method that is closer to how humans behave that yields better results? Some results concerning these questions can be found in our publications (Rickert, 1998; Rickert and Nagel, 1998; Nagel and Barrett, 1997; Nagel et al, 1998c; Kelly and Nagel and Kelly, 1998). The most important ones are maybe that (i) we never encountered a situation where the relaxed state was path dependent, and that (ii) we indeed found relaxation schemes that were significantly faster than others. Note that stochasticity is probably very important for (i).

4 Credibility of results

In order for model results to be useful, they need to be credible. But how can credibility be established? In principle, one should apply the model to a wide range of forecasting scenarios, and build trust over time by demonstrating to which extent the forecasts turn out to be correct. Yet, this approach is made difficult for planning applications because of 20-year forecasting horizons and also because of the stochasticity of the problem.

Because of the 20-year problem, people generally resort to "forecasting" today's traffic using the same input data they would have for a 20-year forecast (see Nagel et al, 1998a, for such a study in the TRANSIMS context). Even then, the stochasticity problem remains, meaning that one would have to generate a *distribution* of scenarios and compare them against a *distribution* of field measurements.

If one finds that the model forecasts deviate from the field measurements, it becomes important to find out the reason for this. The straightforward way to resolve this – to test each individual module of the model package – is again difficult for transportation planning applications since field data for sufficiently controlled situations is hard to obtain. What is thus necessary is to replace modules by other ones and check if results get better or worse.

We have done exactly these studies for the microsimulation module of TRANSIMS (Nagel et al, 1998c, Nagel et al, 1998a). The overall result was that changing the micro-simulation did not lead to huge changes in the results. This indicates that, in first order, deviations from reality were probably not caused by the details of the microsimulation technology, but rather by the demand generation or by the routing. Near-term research in this area should therefore probably focus on the latter.

5 Computing

The current computing requirements for TRANSIMS reach from significant to enormous. For example, a so-called queue model microsimulation (Simon and Nagel, 1999) on the 20 000 link network that Portland (Oregon) uses for EMME/2 studies runs more than 30 times faster than real time on a single Pentium CPU, even with 70 000 vehicles simultaneously in the simulation. Running from 4am to noon would thus take 15 minutes of computer time. However, it is necessary to run this at least 50 times until a solution is reached where both routes and activities are sufficiently relaxed (Esser and Nagel, 1999). Neglecting computer time for re-planning, this results in a computer time of more than 12 hours – feasible, but still significant.

Now, when using a more realistic simulation (the TRANSIMS microsimulation) and a more complete network of 200 000 links, the computing speed drops to four times faster than real time even on a supercomputer with 256 CPUs (Nagel et al, 1998b). If one is interested in a 24-hour period, running the 50 iterations would now take more than 12 days of continuous computing, and remember – this was on a supercomputer.

In conclusion, although it is now possible to run transportation microsimulations of enormously large problems, the computing aspects are still challenging and demand a knowledgeable use of the available technology.

6 Summary

An agent-based simulation approach to transportation is possible with current technology, although it is still a significant computational challenge, especially when one strives for a realistic representation of the outside world. It is possible to break up the problem into modules for "population generation", "activities plans generation", "transportation plans generation", "micro-simulation, and "analysis". The computational challenge stems from the necessity to simulate realistic second-by-second driving of metropolitan regions with many millions of travelers. It is made even harder by the fact that the micro-simulation needs to be run many times so that the plans from the plans generation modules can relax towards "consistency" with the micro-simulation.

7 Acknowledgment

This work was performed at Los Alamos National Laboratory, which is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36 (LA-UR 98-4389).

References

[1] C. L. Barrett, R. Jacob, and M. V. Marathe. Formal language constrained path problems. Los Alamos Unclassified Report (LA-UR) 98-1739, see transims.tsasa.lanl.gov, 1997.

[2] 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.

[3] R.J. Beckman et al. TRANSIMS–Release 1.0 – The Dallas-Fort Worth case study. Los Alamos Unclassified Report (LA-UR) 97-4502, see transims.tsasa.lanl.gov, 1997.

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

[5] J. Esser and K. Nagel. Census-based travel demand generation for transportation simulations. In W. Brilon, F. Huber, M. Schreckenberg, and H. Wallentowitz, editors, *Traffic and Mobility: Simulation – Economics – Environment*, pages 135–148, Aachen, Germany, Sep/Oct 1999.

[6] Forschungsverbund für Verkehr und Umwelt (FVU) NRW. See `www.zpr.uni-koeln.de/Forschungsverbund-Verkehr-NRW`.

[7] R. R. Jacob, M. V. Marathe, and K. Nagel. A computational study of routing algorithms for realistic transportation networks. *ACM Journal of Experimental Algorithms*, in press. See www.inf.ethz.ch/~nagel/papers.

[8] T. Kelly and K. Nagel. Relaxation criteria for iterated traffic simulations. *International Journal of Modern Physics C*, 9(1):113–132, 1998.

[9] S. Krauß. *Microscopic modeling of traffic flow: Investigation of collision free vehicle dynamics.* PhD thesis, University of Cologne, Germany, 1997. See www.zpr.uni-koeln.de.

[10] MITSIM. Massachusetts Institute of Technology, Cambridge, Massachusetts. See its.mit.edu.

[11] K. Nagel. Experiences with iterated traffic microsimulations in Dallas. In D.E. Wolf and M. Schreckenberg, editors, *Traffic and granular flow'97*, pages 199–214. Springer, Heidelberg, 1998.

[12] K. Nagel and C.L. Barrett. Using microsimulation feedback for trip adaptation for realistic traffic in Dallas. *International Journal of Modern Physics C*, 8(3):505–526, 1997.

[13] K. Nagel, M. Pieck, P. M. Simon, and M. Rickert. Comparison between three different micro-simulations and reality in Dallas. Los Alamos Unclassified Report (LA-UR) 98-2944, see www.inf.ethz.ch/~nagel/papers, 1998a.

[14] K. Nagel, M. Rickert, R. Frye, P. Stretz, P. M. Simon, R. Jacob, and C. L. Barrett. Regional transportation simulations. In *Proceedings of the Advanced Simulation Technologies Conference*, Boston, MA, U.S.A., 1998b. The Society for Computer Simulation International.

[15] K. Nagel, M. Rickert, and P. M. Simon. The dynamics of iterated transportation simulations. Presented at the 3rd TRIannual Symposium on Transportation ANalysis (TRISTAN-III) in San Juan, Puerto Rico, 1998.

[16] K. Nagel, P. Stretz, M. Pieck, S. Leckey, R. Donnelly, and C. L. Barrett. TRANSIMS traffic flow characteristics. Los Alamos Unclassified Report (LA-UR) 97-3530, see www.inf.ethz.ch/~nagel/papers, 1997. Earlier version: Transportation Research Board Annual Meeting paper 981332.

[17] M. Rickert. *Traffic simulation on distributed memory computers.* PhD thesis, University of Cologne, Germany, 1998. See www.zpr.uni-koeln.de/~mr/dissertation.

[18] M. Rickert and K. Nagel. Issues of simulation-based route assignment. Presented at the International Symposium on Traffic and Transportation Theory (ISTTT) in Jerusalem, 1999.

[19] P. M. Simon and K. Nagel. Simple queueing model applied to the city of Portland. *International Journal of Modern Physics C*, 10(5):941–960, 1999.

[20] TRANSIMS, TRansportation ANalysis and SIMulation System, since 1992. See transims.tsasa.lanl.gov.

[21] M. D. Williams, G. Thayer, and L. L. Smith. A comparison of emissions estimated in the transims approach with those estimated from continuous speeds and accelerations. Los Alamos Unclassified Report (LA-UR) 98-3162, see transims.tsasa.lanl.gov, 1998.

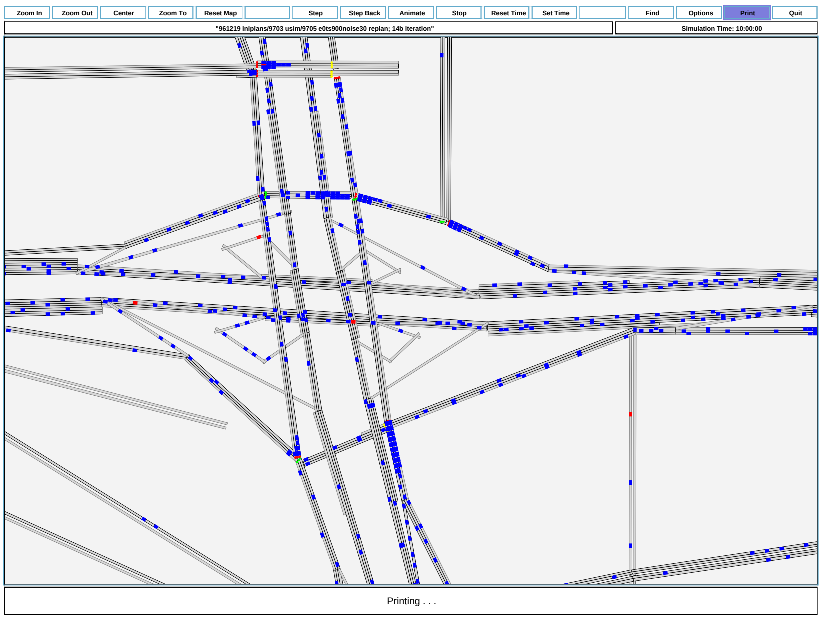

Figure 1: Zoomed-in snapshot of a micro-simulation run.

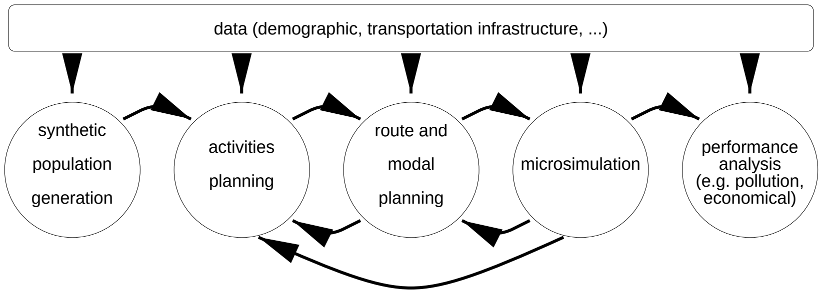

Figure 2: Modules of an agent-based simulation approach to transportation.

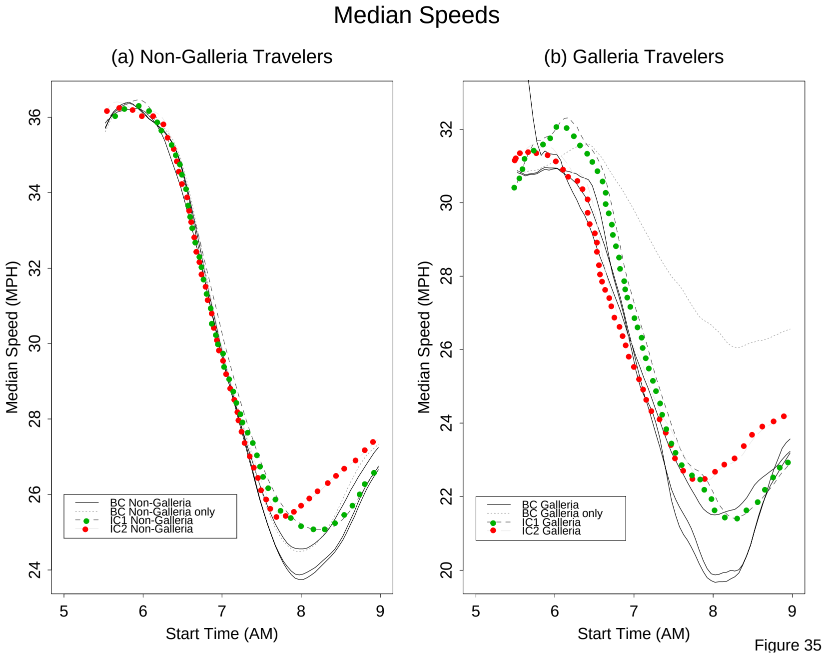

Figure 3: Example output of a case study. The curves show an average speed in a study area as a function of time. Clearly, speeds start out high, then drop during the rush period, and recover afterwards. Different lines denote results for different transportation infrastructures that were compared in the study. From Beckman et al, 1997.

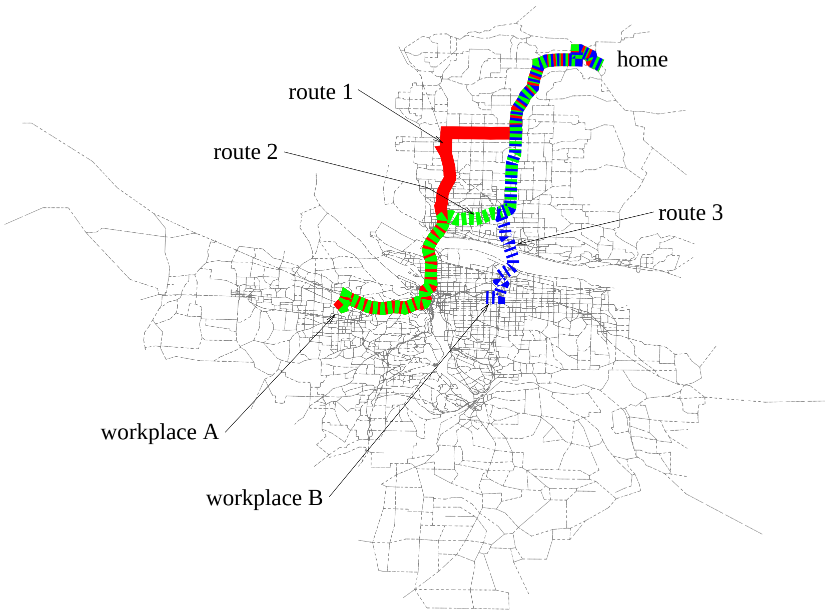

Figure 4: Example of re-planning. The traveler first drives to workplace A via route 1. Eventually, he/she decides to switch to a supposedly better route 2. Again later, he/she decides that this results still in too much driving time, and moves to a different workplace B closer to home. The street network shows Portland/Oregon. The example is taken from an actual feedback series (Esser and Nagel, 1999).