LA-UR 98-312
Approved for public release; distribution is unlimited

Regional Transportation Simulations

Kai Nagel, Marcus Rickert, Roger Frye, Paula Stretz, Patrice Simon, Riko Jacob, Christopher L. Barrett

Published in: High Performance Computing 1998, edited by Adrian Tentner, The Society for Computer Simulation International, San Diego, CA, U.S.A., p 104–109.

LOS ALAMOS
NATIONAL LABORATORY
Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

REGIONAL TRANSPORTATION SIMULATIONS

Kai Nagel, Marcus Rickert, Roger Frye, Paula Stretz, Patrice Simon, Riko Jacob, Christopher L. Barrett

Los Alamos National Laboratory
TSA-DO/SA, Mail Stop M997
Los Alamos NM 87545, U.S.A.
kai@lanl.gov

Keywords: transportation simulation, TRANSIMS, simulation of socio-economic systems

Abstract. For transportation planning applications, it is useful to not only model each individual traveler, but also the decision-making process leading to her travel demand. Simulation-based modeling of this process means iterations between the actual transportation micro-simulation and the modules simulating the process making the plans. This means that for understanding a single day of travel, it may be necessary to simulate that day hundreds of times for the iteration process, leading to a considerable strain on computational resources.

1 Introduction

Transportation forecasting has a short-term and a long-term component. A typical short-term forecast may concern the traffic half an hour into the future. This traffic may be different from the day before because of weather, an accident, or simply the stochasticity of the system. Long-term forecasts typically are transportation *planning* questions, for example of how the traffic will be in twenty years from now given certain demographics and a certain transportation infrastructure. Transportation forecasting shares this distinction between short-term and long-term forecasting with the atmospheric sciences. In both areas, methods for both questions can be similar, but the focus of the models usually is different. For example, for short-term forecasting, it does not matter if travelers are going to give up a trip the next day because of a new bottleneck (say, a construction site), whereas for long-term forecasting, the answer to such a question is crucial since it will describe the reaction of travelers to infrastructure changes. Transportation science also shares with the atmospheric sciences that large scale controlled experiments are impossible since nobody can afford to change the transportation infrastructure of a whole region just for experimental purposes.

Yet, there are also significant differences between the two. The most important one probably is that humans do not only react to "short range" forces: If they expect a traffic jam tomorrow, they may do the shopping already today; and if one makes a forecast publicly available, it will probably render itself irrelevant because people react to it. This means that a "one-shot" approach would not describe reality; indeed, reality itself is iterative, with a complicated interaction between anticipation and reaction. As an example, let a long-expected bottleneck appear in a transportation system. Some people will have adapted to it in anticipation. Yet, the traffic pattern that develops will be different from the anticipated pattern. So people react to the actually occuring pattern, test their "strategies" again the next day, adapt again, etc.

This implies that transportation simulations on that scale need to incorporate iterations. It is these iterations which represent a considerable part of the current computational and methodological challenges of transportation simulations: computationally, because for a forecast, one needs to run the scenario over and over again, allowing for adaption; methodologically, because it is close to impossible to simulate what people really do and so one has to hope that surrogate methods will lead to real enough results to be realistic. A simple but unrealistic reactive strategy to deal with the bottleneck situation may lead to the same overall traffic pattern, but may need more iterations than reality.

This paper will concentrate on long-term transportation forecasting. It will outline how one current project in that area, the TRANSIMS (TRansportation ANalysis and SIMulation System) project at Los Alamos National Laboratory, approaches the problem, how much of this is already implemented and tested, and what the implications for computing needs are.

2 TRANSIMS

The probably "cleanest" approach to a problem such as traffic forecasting is a "microscopic" approach, i.e. an approach where each entity of the problem is represented individually. For traffic, that currently means that individual travelers need to be represented. Since

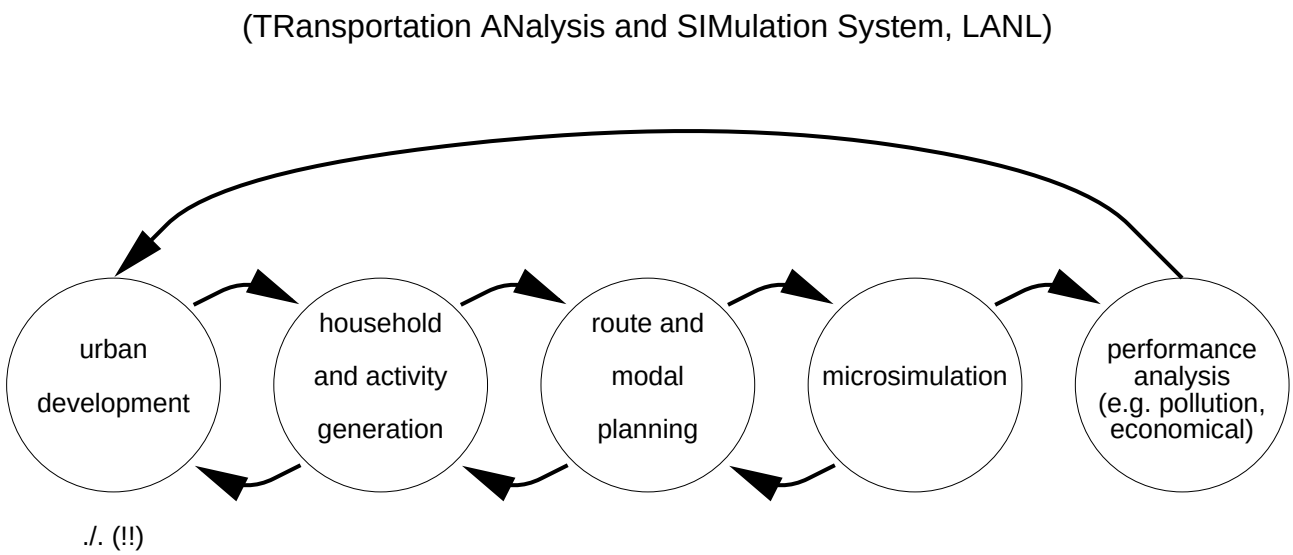

Fig. 1. TRANSIMS design. Urban development is integral part of the causal loop, but is not included in the current TRANSIMS project.

analytical methods on that level are often unable to handle the complexity of real world problems, simulation is an attractive option.

The TRANSIMS (TRansportation ANalysis and SIMulation System) project at Los Alamos National Laboratories is such a "microscopic" transportation simulation project. It attempts to simulate all aspects of human behavior that are relevant to transportation planning. The design (see Fig. 1) starts from demographic data and, as a first step, generates synthetic individuals and households plus their activities (such as work, shop, be at home, etc.). Since all the activities are localized on the street network, this induces travel demand. From here, the synthetic individuals make modal and route choices, and finally, all "plans" are executed in a microscopic simulation of all travelers. The simulation result can be used as the basis for further analysis, such as stake-holder analysis or emissions modeling.

Note that everything in this design operates on the level of individual travelers. Although one certainly has to accept simplifications in the logic of human decision-making, this allows at least in principle to identify processes that occur in the real world, and to make the model arbitrarily complex if this is desired for a certain problem.

As outlined above, there is no hope of making the causality uni-directional. Congestion showing up in the micro-simulation will cause people to re-route, to choose a different mode, to change the locations of their activities or their activities in general, or even to move their home to a different location. For that reason, the TRANSIMS design allows for feedback between the modules. It is clear that this puts an enormous load on the computational demands, since the microsimulation has to be run many times until a "relaxed" result is obtained.

3 Route re-planning

Currently implemented in TRANSIMS is the feedback between routing and micro-simulation; the feedback into higher levels of the planning process (location choice, activities planning) is under development. The feedback has been tested using a situation in the Dallas/Fort Worth area as a study case. This is what has been done:

- Since digital network data of the whole area was not available, a so-called "focused" network was used that contained all streets in a 5 miles $\times$ 5 miles "study area", but was considerable "thinned out" with further distance from there. The complete network had 14751 links and 9864 nodes.
- Since the activities generation part of TRANSIMS is not yet operational, trip tables were used as an interim method. The Dallas trip tables contain 24-hour counts of trips between different zones in the area. One of the problems with trip tables is that their information is routinely unreliable (how

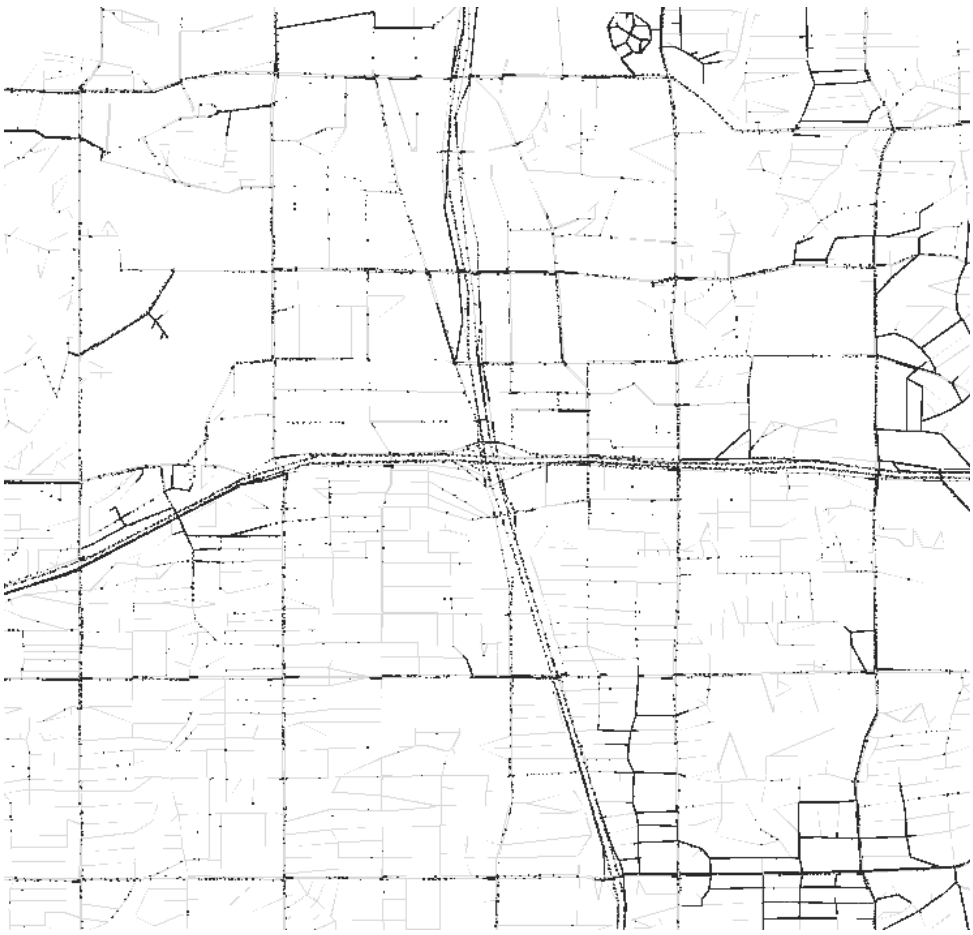

Fig. 2. "Initial" traffic. Note the excessive jams in the residential streets.

do you "measure" the required information?), and they are wrong after major infrastructure changes, making them useless for transportation planning purposes. Yet, they seem useful enough for an interim study.

- The 24-hour trip tables were converted into trip tables that reflected the time-of-day, for example the large amounts of home-to-work trips during the morning rush hour [1, 2]. After that, the tables were converted into lists, where each entry in the list contains a trip, i.e. a starting time, a starting location, and a destination location. There were about 10 million such trips during the 24-hour period in the Dallas/Fort Worth area (population approx. 3.5 million).
- Out of these 10 million trips, all of them starting between 5am and 10am (approx. 3 million) were routed through the network, using fastest path in the empty network ("initial routing"). Obviously, this does not take congestion into consideration. Only the trips going through the study area were retained after this, approx. 300 000 trips.
- These 300 000 trips were run through a micro-simulation. Note that this implies that the micro-simulation executes pre-computed routes. The micro-simulation records travel times through links as a function of time-of-day; obviously, congested links will report long link travel times. The micro-simulations only simulated a 5 miles × 5 miles study area, which has 6124 links and 2292 nodes, with overall 2276 lane-kilometers.
- The link travel times are used to compute new routes for a fraction of all travelers. The intuition behind this is that a certain number of travelers decides, over night, to try a different route the next day; the information they have available for this are all link travel times from the previous day (an unrealistic assumption in real-world behavioral terms).
- The last two steps are run over and over again until some "relaxation" is found. When using stochastic micro-simulations (as we do), "relaxation" is not well defined, so it is an area of research [8, 3, 4].

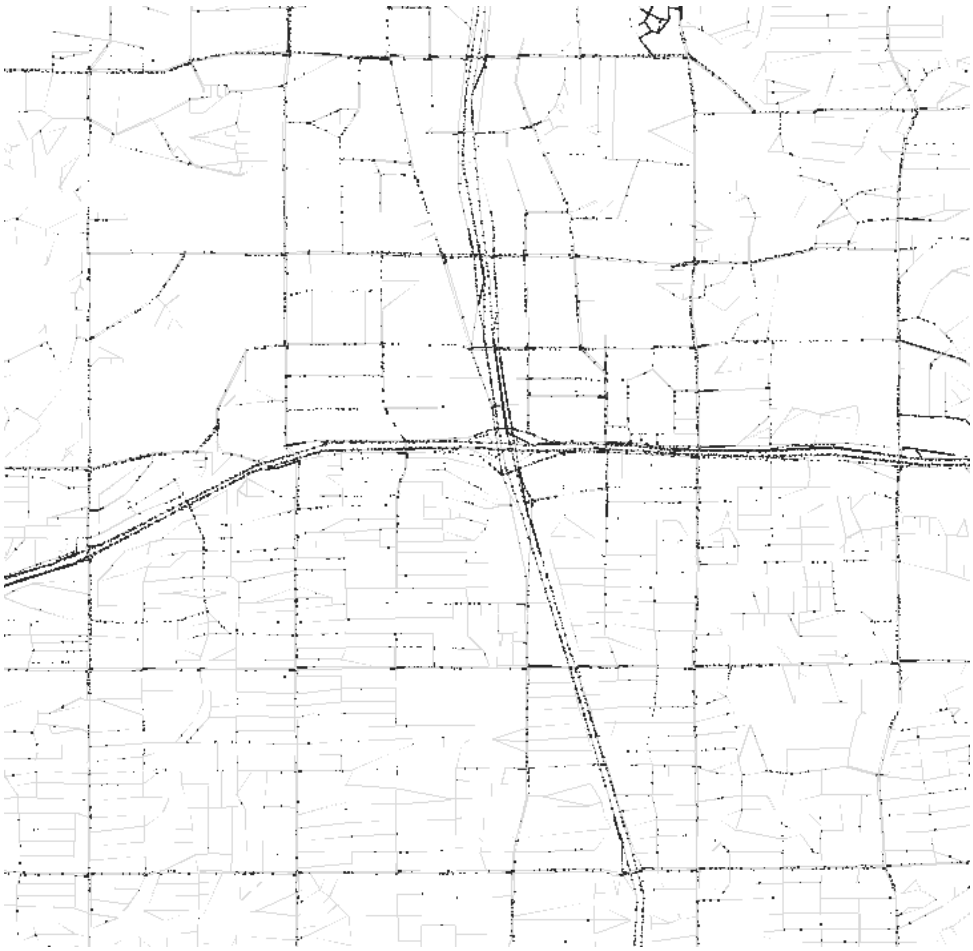

Fig. 3. "Relaxed" traffic

Figs. 2 and 3 are examples of "initial" and "relaxed" traffic situations. Clearly, "relaxation" means that traffic is better distributed across the system and jams have dissolved. Further details can be found in [6, 3, 8, 4, 5].

4 Computation

We currently have three micro-simulations, which differ in the amount of realism. The most realistic one [7] runs, for the given problem, about as fast as reality on five coupled SUN Sparc5. The second, less realistic one (not having turn pockets, having only "average" traffic lights, etc.), called PAMINA [9, 3, 8], runs about 20 times faster than real time on 6 CPUs (250 MHz) of a SUN Enterprise 4000. Both simulations use PVM for message passing; PAMINA can also use MPI. The last, again less realistic micro-simulation, runs about 20 times faster than real time on a single 250 MHz CPU [10]. It seems that, given current computing and communications technology and using a bus technology (Local Area Network or Shared Memory Computer) for communications, 20 times faster than reality is a rough upper bound on computing speed [8]. Using a two-dimensional communications technology gets around this bottleneck.

Let us, for the sake of the argument, focus on PAMINA which is also the one that is best tested in terms of computational performance. For systematic relaxation studies using different iteration schemes, about 1000 runs of the morning (5am to noon) were run. This means 350 hours of continuous computation on the above-mentioned 6 CPUs of the Enterprise. A major result of these studies was that, with certain relaxation schemes, 20 iterations can be enough to reach a "relaxed" traffic state [8, 3]. – Further computational time is needed for the routing calculations. We can currently compute about 300 routes per second per 250 MHz CPU for the given network size; that is, the initial routing run (3 million plans) using 6 CPUs needs about 1/2 hour; re-planning runs then need an amount of time that is negligible compared to the micro-simulation be-

cause only 300 000 plans are left and only a fraction of those is re-planned in each iteration [?].

Note that all these studies were run on a relatively small portion of Dallas. Using study-areas like this causes many problems. For example, the re-planner re-routes trips around the study area because congestion only happens inside the simulated study area. Or it is really hard to get reasonable "measures of efficiency" for most of the trips since a large portion of the trip can be outside the study area and thus has unreliable travel times. In general, it seems desirable to run simulations on complete metropolitan areas until the effect of artificial boundaries is sufficiently studied. TRANSIMS will, as its next study, investigate traffic in the Portland area. Portland is still a moderately small problem, with 1.6 million people living in the metropolitan area. The road network for the whole area has 120 000 links, considerably more than the focused Dallas network.

In order to plan all trips in the area, this would mean about 2 hours of computing time on our 14-CPU machine; re-planning runs would be a factor of ten or more shorter because only a fraction of the plans is re-routed in each iteration.

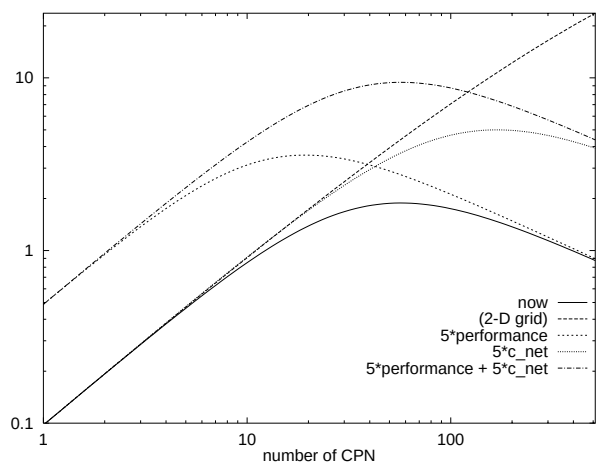

Fig. 4. Expected computing speed (real-time ratio) for Portland road network. Shown are: current situation (250 MHz CPU and bus communication with FDDI characteristics); same CPUs but 2-D grid communications topology; CPUs that are 5 times faster but unchanged communications technology; faster CPUs and faster communications.

In order to obtain an estimate for the microsimulation, one needs to extrapolate the computing times for the approx. 6000 link Dallas study area network to the 120 000 link Portland network. The result of this extrapolation is shown in Fig. 4 (for the method see [8]). One can see that, given our current technology, the simulation would be as fast as reality (real-time ratio RTR = 1) on 12 CPUs; adding CPUs could push it to twice as fast at 55 CPUs, but not beyond this. Using five times faster CPUs would result in a maximally achievable RTR of 4 with 20 CPUs; using five times faster communication but the "slow" CPUs allows for an RTR of 5 with 160 CPUs. Only a changed communications topology really allows to go to really high RTRs using a massively parallel machine even given current technology. – On our own machine, a single 24 hour simulation (as we are planning) would thus take 24 hours. 20 iterations for relaxation, as indicated above, would need at least 240 hours.

On top of this will be the iterations for the activities: People do not only adapt their routes in reaction to congestion, but also the locations of their activities, their sequencing, or they will give up some of them completely. Let us assume that we need overall 50 runs of the microsimulation, and that no additional computational load will occur due to the activities scheduling problem. That means that a single study will still take 1200 hours, or 50 days, or nearly two months, of pure computing time (i.e. not counting hardware failures, disk space problems, etc.). Note that these times will become considerably worse with our more realistic microsimulation. It is possible that for many problems even less realistic micro-simulations turn out to be sufficient; but systematic research comparing to using more realistic micro-simulations will be necessary.

In conclusion, it seems that it will be possible to run micro-simulation approaches to transportation planning of regional areas in the near future on multi-CPU desktop workstations, although patience by the user will be required. It is also clear that for research purposes and for larger systems, there will still be a long time into the future where more massively parallel approaches would pay off, even for the problem sizes that we are attacking today.

5 Summary

The probably most systematic approach to transportation forecasting are microsimulations, i.e. simulations where each individual car is resolved. Using microsimulations for transportation planning purposes, i.e. for forecasts twenty years or so into the future, means that the human decision-making process related to the planning of transportation needs to be included in the model. This process is, in spite of all human intelligence, iterative. For computational models, one can to a certain extent replace human intelligence by doing even more iterations. Rough estimates for the necessary computer time for such an iteration process show that it will probably be possible to run such studies on powerful multi-CPU desktop workstations, but the

user has to be patient enough to live with turn-around times of a month (!), although research results may point to methods to make this time shorter. For research purposes, these computing times are most probably unacceptable, and massively parallel approaches should be used.

References

1. R.J. Beckman. Personal communication.
2. R.J. Beckman et al. TRANSIMS–Release 1.0 – The Dallas-Fort Worth case study. Technical Report 97-4502, 1997.
3. H.J. Herrmann, E. Flekkoy, K. Nagel, G. Peng, and G. Ristov. Density waves in granular flow, 1994.
4. T. Kelly and K. Nagel. Relaxation criteria for iterated traffic simulations. *International Journal of Modern Physics C*, 9(1):113–132, 1998.
5. K. Nagel. Experiences with iterated traffic microsimulations in Dallas. pages 199–214.
6. 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.
7. K. Nagel, P. Stretz, M. Pieck, S. Leckey, R. Donnelly, and C.L. Barrett. TRANSIMS traffic flow characteristics. Technical Report 97-3530, 1997. Earlier version: Transportation Research Board (TRB) preprint 981332.
8. M. Rickert. *Traffic simulation on distributed memory computers*. PhD thesis, University of Cologne, Cologne, Germany, 1998.
9. M. Rickert and K. Nagel. Experiences with a simplified microsimulation for the Dallas/Fort Worth area. *International Journal of Modern Physics C*, 8(3):483–504, 1997.
10. P.M. Simon and K. Nagel. A simplified cellular automaton model for city traffic. *Physical Review E*, 58(2):1286–1295, 1998.