Los Alamos Unclassified Report LA-UR 98-2944 (This version: March 24, 1999)
Approved for public release; distribution is unlimited

Comparison between three different traffic micro-simulations and reality in Dallas (comments welcome)

Kai Nagel, Martin Pieck, Patrice M. Simon, Marcus Rickert

LOS ALAMOS
NATIONAL LABORATORY

This page intentionally left blank.

Comparison between three different traffic micro-simulations and reality in Dallas (comments welcome)

Kai Nagel,[a,b] (505) 665-0921 phone, (505) 982-0565 fax, kai@lanl.gov

Martin Pieck,[a] (505) 665-0086 phone, pieck@lanl.gov

Patrice M. Simon,[a,b] (505) 667-6143 phone, simonp@lanl.gov

Marcus Rickert[a,b] (505) 665-6608 phone, rickert@lanl.gov

[a] Los Alamos National Laboratory, TSA-DO/SA, Mail Stop M997
P.O.Box 1663, Los Alamos, NM 87545, USA

[b] Santa Fe Institute, 1399 Hyde Park Rd, Santa Fe NM 87501, USA

March 24, 1999

Abstract

We describe three traffic microsimulations which operate at different levels of fidelity. They are used to iteratively generate a self-consistent route-set based upon microsimulation feedback. We compare the simulation results of all three simulations to aggregated turn count data of actual field measurements.

Acknowledgments – We thank A. Bachem and R. Schrader for supporting MR's work as part of the traffic simulation efforts in Cologne ("Forschungsverbund Verkehrsimulation und Umweltwirkungen NRW" [1]). We also thank them, plus C. Barrett, P. Wagner, Ch. Gawron, S. Krauß, and M. Schreckenberg for help and discussions. We thank the TRANSIMS research (M. Marathe, R. Jacob, and others) and software (M. Stein, P. Stretz, D. Roberts, B. Bush, K. Bergbigler, and others) teams for discussions and the processing of the data which served as input for the micro simulation.. The North Central Texas Council of Government (NCTCOG/Ken Cervenka) provided road network and origin-destination data on which the simulations were based. Use of computer facilities at the Center of Parallel Computing (Cologne) and computing time on the workstation cluster at TSA-DO/SA is gratefully acknowledged. We further thank all people in charge of maintaining the above mentioned clusters. The work of MR was supported in part by the "Graduiertenkolleg Scientific Computing Köln". This work was performed in part 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 98-2944).

1 Introduction

It is certainly desirable that transportation forecasting models are "correct" in the sense that the traffic patterns they predict correspond to what would happen in reality under the circumstances assumed in the forecasting model. Unfortunately, it is notoriously difficult to transform the above common sense statement into a technical specification. Since one cannot run controlled experiments in socio-economic systems, it is usually impossible to check the forecasts. Let us mention some of the problems:

- It is very difficult to obtain "clean" field measurements of traffic flow characteristics such as flow–density curves. The problem is that one needs crowded situations, and crowded regions usually have plenty of "messiness" such as on-/off-ramps, intersections, weaving sections, etc. which influence the measured quantity and thus make it dependent on the dynamical context. For example, a downstream bottleneck will cut off most of the data near capacity for a flow-density diagram.

- It seems thus that one should add the "context" to the simulation model, i.e. include the complexity of the real world around the measurement site. The problem now is that one suddenly needs more information about the vehicles, such as their routes and intended movements through the system. For example, it clearly is important how much of the traffic streams in a weaving section is crossing to the other side[2].

- Yet, reliable origin-destination information is difficult to obtain. In theory, it would be possible to collect for the above situation enough information using license-plate detection. In practice, this is usually costly.

- A possible way out is to generate the origin-destination information from "activities", that is to simulate the complete human decision-making process related to transportation, starting from defining the activities (be at home, work, eat, shop, ...) and the activity locations [3, 4], via selecting mode and routes, to driving decisions such as lane changing and gap acceptance. The consistent simulation of this approach in a unified framework is, in very short, the idea of the TRANSIMS (TRansportation ANalysis and SIMulation System) project [5].

The net result is that it is very hard to compare simulation results to reality. For example, if a simulation is driven by route plans (as in this paper), are differences to field data in the intersection turn counts due to wrong origin-destination tables, due to wrong routing, or due to wrong traffic flow dynamics?

In our intuition, there is currently no satisfying way out of the dilemma (and maybe there will never be). Yet, it is certainly possible to do systematic studies. For example, the sensitivity of turn counts on variations of the origin-destination tables can be tested. Or one can document emergent traffic flow behavior for simplified, "clean" cases (such as saturation flow from a minor into a major road) [6]. Comparison of these cases with reality will still be difficult, but at least one can compare simulations with each other. In this way, it should be possible to systematically enhance our understanding of the intricacies of the simulated dynamics and maybe gain enough practical experience to also be able to say something about the forecasting quality.

This paper contributes a piece to the mosaic. It describes micro-simulations that have been made using data for the Dallas/Fort Worth area (described in Sec. 2). The micro-simulations (described in Sec. 3) run on routes; the routes are generated iteratively using fastest path (see Sec. 4). For comparison purposes, we had turn counts for specific intersections for the morning peak (described in Sec. 5) and results of an earlier traditional assignment for the same region (described in Sec. 6). We do comparisons using approach volumes (Sec. 7) and turn counts (Sec. 8). The paper is concluded by a discussion (Sec. 9) and a summary (Sec. 10).

2 Context

The context of this paper is the so-called Dallas–Fort Worth case study of the TRANSIMS project [7]. Most of the details relevant for this paper can also be found in Ref. [8]. Purpose of the case study was to show that a micro-simulation based approach to transportation planning such as promoted by TRANSIMS will allow analysis that is difficult or impossible with traditional assignment, such as measures of effectiveness (MOE) by sub-populations (stake-holder analysis), in a straightforward way. Most of the accompanying studies such as Refs. [9, 8, 6, 10, 11] and also this paper attempt to document the technology leading to and following up on the case study.

The underlying road network for the study (public transit was not considered) was a so-called focused network. It contained *all* links in a 5 miles times 5 miles study area, but got considerably "thinner" with further distance from the study area.[1] A picture of the focused network can be found in Ref. [8].

The TRANSIMS design specifies to use demographic data as input and generate, via synthetic households and synthetic activities, the transportation demand. The Dallas/Fort Worth case study was based on interim technology: parts of the demand generation modules were not yet available. For that reason, the study uses conventional 24-hour trip tables (production-attraction matrix, PA matrix) as starting point. The PA matrix was provided by the regional transportation planning authority (the North Central Council of Government, NCTCOG). The PA matrix roughly is a 24 hour origin-destination matrix, i.e. the metropolitan area of Dallas/Fort Worth is divided into approximately 800 zones (traffic analysis zones, TAZs), and the number of trips going from each zone to each other zone in a 24 hour period is given.

For the case study, the first thing that was done was to break down the PA matrix into individual trips [7]. For this, a time-of-day distribution according to land use in the destination zone was used. For example, traffic going to commercial zones mostly occurs in the morning. Also, starting and ending locations of trips were specified on the link level. The result was a table of approx. 10 million trips, all with a starting time, a starting location, and a destination location. From this table, all trips starting between 5am and 10am (ca. 3 million trips) were actually used.

Next, an "initial routing" step was done. It is easiest to imagine that all trips were routed according to "fastest path in an empty network" (i.e. using free speeds provided by the transportation authority). All trips that in this step went through the study area were retained, all

[1] Note that this "thinning out" of the network was not done in any systematic way and is explicitly *not* recommended. It was an ad-hoc solution because more data was not available, and because of limited computing capabilities.

other trips were removed. Note that this defines a base set of trips for all subsequent studies presented in this paper: All trips thrown out in this step can no longer influence the result of the studies. This base set contained approx. 300 000 trips.

For the results in this paper, two different base sets of trips were used. That is, the initial routing for the case study was not done using fastest path in an empty network, but instead some untested and undocumented variation of an assignment technique. In essence, it routed some trips, calculated new link travel times based on a standard link performance function, routed more trips, etc., until all trips were routed. Since it is somewhat unpredictable what this exactly does, later research studies were based on a base set of trips obtained by routing in an empty network.[2]

3 The micro-simulations

The above procedure does not only generate a base set of trips, but also an initial set of routes (called *initial planset*). These routes are then run through a traffic micro-simulation, where each individual route plan is executed subject to the constraints posed by the traffic system (e.g. signals) and by other vehicles. Note that this implies that the micro-simulation is capable of executing pre-computed routes (only very few micro-simulation currently have this capability although their number is growing), and it also implies that, in the simulations, drivers do *not* have the capability of changing their routing on-line.[3] Three micro-simulations have been used, all three related to the TRANSIMS project, but with different levels of realism and different intended usages. For simplicity, we will just number them, i.e. "micro-simulation 1", "micro-simulation 2", and "micro-simulation 3" (MS1, MS2, and MS3). MS1 is the most realistic one, MS3 the least realistic one of the three. MS1 and MS2 are based on the so-called cellular automata technique for traffic flow [13, 14, 15] although there is no necessity for this except the requirement of sufficient computational speed. MS3 is based on a simple queueing model in the spirit of previous work found in [16, 17] and [18]; see [**?**] for more information.

3.1 Micro-simulation 1 (MS1, TRANSIMS)

MS1 is the "mainstream" TRANSIMS micro-simulation. As said above, it is the most realistic of the three, including elements such as number of lanes, speed limits, (fixed) signal plans, weaving and turn pockets, lane changing both for vehicle speed optimization and for plan following, etc. It also has the most sophisticated output subsystem of the three, allowing the user to specify which data to collect during the simulation. The studies described on this paper were run on five coupled Sparc5 workstations running as fast as real time; newer versions also run on an Enterprise 4000. Details of this micro-simulation are documented in [6, 19].

[2] As already stated above, we explicitly do not recommend using a study area as we did for Dallas because of a large number of currently unsolved associated problems. When simulating the whole metropolitan area, the initial routing should not matter for the "relaxed" result and so the problem would go away.

[3] It is not that on-line re-routing is incompatible with the technology (see, e.g., [12]), but it has not generally been implemented and studied.

3.2 Micro-simulation 2 (MS2, PAMINA)

The second micro-simulation, MS2, does not include signal plans, weaving and turn pockets, and lane changing for plan following. Most other specifications are the same as for MS1, although further differences can be caused by the different implementation. MS2 is much better optimized for high computing speed: it ran more than 20 times faster than MS1 for this study, which is a combined effect of using faster hardware (the code is much easier to port to different hardware, thus being able to take advantage of new and faster hardware much sooner), less realism, and an implementation oriented towards computational speed. This micro-simulation is documented in [20, 21, 10, 12].

3.3 Micro-simulation 3 (MS3, QM)

The third micro-simulation that we used is significantly less realistic than the other two. In this model, each link is represented by a queue with a service rate proportional to its capacity. In addition, a link has a limited "storage capacity", representing the number of vehicles that can sit on the link at jam density.

When a car enters a link at time t_{enter}, an expected link travel time, T_{free}, is calculated using the length and the free flow speed of the link. The vehicle is then put into the queue, together with a time $t_{earliest} = t_{enter} + T_{free}$ which marks the earliest possible departure at the other end of the link. In each time step (which we take as one second), the queue is checked if the first vehicle can leave according to $t_{earliest}$, according to the capacity constraints, and according to the storage constraints of the destination link. The queue is served until one of these constraints is not fulfilled.

Note that we do not use a fundamental diagram 'velocity versus density' to compute the travel time, as proposed in [16]. Indeed, our simple queueing model generates a reasonable fundamental diagram without any further input. More details can be found in [**?**].

The reason for having a model like this is that we want a micro-simulation model that fits into the overall TRANSIMS framework (i.e. runs on individual, pre-computed plans) but has much less computational and data requirements than the other simulation models. Indeed, MS3 runs on the same data as traditional assignment models, and on a single CPU it is computationally a factor 60 faster than MS2. A parallel version is planned.

4 Feedback iterations and re-planning

The initial plan-set is obviously wrong during heavy traffic because drivers have not adjusted to the occurrence of congestion. In reality, drivers avoid heavily congested regions. We model that behavior by using *iterative re-planning* [22, 23, 24]: The micro-simulation is run on a pre-computed plan-set, and travel times along links are collected. Then, for a certain fraction, X, of the drivers, new routes are computed based on these link travel times. Technically, each route from the old plan-set is read in, with probability $1 - X$ is is written unchanged into a new file, and with probability X a new route is computed given the starting time, starting location, and destination location from the old route plus the time-dependent link travel times provided by the last iteration of the micro-simulation. After this, the micro-simulation is run again on the new plan-set, more drivers are re-routed, etc., until the system is "relaxed", i.e.

no further changes are observed from one iteration to the next except for fluctuations (all micro-simulations are stochastic).

We have used two different implementations of the re-planner. For technical completeness, let us call them RP1 and RP2. RP1 is the re-planner that was used for the Dallas/Fort Worth case study; in this paper, it is used in conjunction with micro-simulation MS1. RP2 is a faster and less memory-consuming version that has been implemented since then. RP1 and RP2 are written according to the same specifications: they compute fastest paths based on 15-minute averages of link travel times using a time-dependent implementation of the Dijkstra algorithm [8, **?**]. Time-dependence is accounted for in the following way: The micro-simulation reports the average link travel time of all vehicles *leaving* the link between, say, 8:00 and 8:15. RP2 then uses this link travel time for all Dijkstra calculations that *enter* the link during the same time period. RP1 uses this link travel time for all Dijkstra calculations that enter the link between 7:45 and 8:00 (thus "anticipating" congestion build-up). Clearly, both algorithms are somewhat sloppy here; newer implementations of our algorithm deal with this in a more precise way by actually calculating when, in the average, the vehicles had entered the link. Both RP1 and RP2 were tested together with the micro-simulation MS2 and no significant differences were seen.

Certainly, there are many questions. How can one tell that an iteration series is relaxed? Do different initial conditions and/or iteration schemes lead to the same overall relaxed state? If so, can one speed up the relaxation process? Is reality at all similar to the relaxed state obtained with this methodology? These questions are treated in other publications [12, 11, 8, 25, 26]. In summary, we can say the following:

- MS1 was iterated using a "scheduled" re-planning fraction, i.e. using a re-planning fraction of 10% for the first seven iterations, followed by five iterations of 5% and two iterations of 2%. Later tests using the same re-planner but a different micro-simulation indicate that the resulting state is not yet relaxed, i.e. one would need to make further iterations to bring, for example, the sum of all travel times to a stable value. See [12, 11, 8] for more details.

- In contrast, MS2 and MS3 were iterated with an "age-dependent" re-planning scheme: The probability of a route being selected for re-planning was made proportional to the number of iterations since the last re-planning event for that route. This was shown to be a much more efficient *and* robust re-planning scheme as any other scheme we tested. See [12, 11] for more details.

- We could not find an indication that the relaxed states depend either on the initial conditions or on the selected relaxation schedule. If this holds in more general, this would be good news because the final state of the iterations would be fairly robust against changes. See [12, 11] for more details.

- However, we could find large fluctuations in the traffic patterns by just changing the random seed, i.e. keeping *everything* (e.g. starting times, individual routes, signal plans) unchanged *except* for changing the sequence of random events that influence acceleration, braking, and lane changing. These fluctuations seem to be non-Gaussian on the system level, i.e. the traffic evolved according to one general pattern most of the time but could be totally different (much more congested) in few of the runs. This is what will be meant by a "typical" vs. a "non-typical" run in the figures. See [25] for more details.

5 Field measurements of turn counts

The regional transportation authority, the North Central Texas Council of Covernments (NCTCOG), performed systematic turn counts on some of the intersections in the study area in 1996. In general, counts are available for through movements, left, and right turns, for the 1-hour period between 8am and 9am. These counts were all done only once, so no information about variability is available. Also, counts for different intersections were done on different days. The locations of the intersections will become clear from the figures presented later; the geographical area in these figures is about a sixteenth of the geographical area that was actually simulated (see, e.g., [8]).

6 NCTCOG 1990 LBJ study results

We also had the results of an 1990 "LBJ corridor study" performed by NCTCOG. The LBJ (Lyndon B. Johnson) freeway is the freeway going in east-west direction through the study area; it can be seen as a gray line near the bottom of the figures shown in this paper. For further information on this study, see [27]. The values that were used for comparison were the 8-9am peak hourly flows.

7 Comparisons using approach volumes

Before turning to quantitative comparisons, it is necessary to point out a consistency problem with the data that significantly reduces comparability. This problem is that the Dallas North Tollway (which is the north-south freeway in the study area; the gray line near the left side of the figures shown in this paper) got extended from the large intersection in the center of the study area (bottom left in the figures in this paper) towards the north sometime around 1990/1991. The result is that the NCTCOG study, our simulations, and the 1996 field data are all inconsistent with each other:

- The NCTCOG study is based on pre-extension trip tables and on a pre-extension road network.
- 1996 reality is based on post-extension trip tables and on a post-extension road network.
- Our simulations are based on pre-extension trip tables but a post-extension road network.

The addition of the freeway most certainly reduced impedance in the north-south direction, thus most probably causing more trips in that direction. This additional travel will be reflected in the 1996 counts, but not in our simulations.

Fig. 1 shows the comparison between the NCTCOG 1990 am peak assignment and the 1996 approach volumes. Wide gray bars show the model results, whereas narrow black bars denote field measurements. The results show a structure that is probably familiar to experts working in the area of assignment:

- At the intersection between Alpha Rd and Preston Rd, in the top right corner of the figure, the southbound approach volume is represented correctly by the model. For most other intersections, model results deviate from the observations, sometimes more, sometimes less. – It seems to be common practice to adjust the assignment models until they get flows on major arterials about right and then accept the results on the other roads. One can thus speculate that the southbound approach to the interesection between Preston and Alpha was one of the "benchmark flows".

- A peculiarity is the northbound approach to the intersection between Alpha Rd and the frontage road in the top left corner of the figure. Here, the assignment exceeds the observation by a factor of four, leading to an hourly assignment of approximately 7000 veh per hour, clearly unrealistic for a 3-lane signalized intersection. The reason for this is that, at the time of the assignment, the Dallas North Tollway freeway, shown as a gray line on the left side of the picture, finished here and ended in the described intersection. Assignment models, in conjunction with the link delay functions commonly used in the States, are known to overassign traffic to short bottlenecks.

In summary, one can say two things about this comparison between 1990 assignment and 1996 observations: (i) Without being able to prove that this is indeed what has happened, the assignment shows structure that can be expected from the technology behind the model (overassignment on short bottleneckts) and from the comman usage of the model (correct assignments on major arterials, deviations everywhere else). (ii) When judging the results, one should, throughout this paper, keep the data inconsistencies in mind.

Figs. 2 (a)–(d) show the same type of comparison, now between the TRANSIMS simulation results and the Dallas observations. Remember that these simulations use the same demand structure (PA matrix) as the NCTCOG assignment but a different network (Dallas North tollway extended towards north). In spite of the huge differences between the three simulation approaches, their approach volume results look remarkably alike. The TRANSIMS 'typical run' (Fig. 2 (a)) and TRANSIMS 'non-typical run' (Fig. 2 (b)) differ mainly on the LBJ west-bound frontage road. This can be explained by the presence of jams for the non-typical run in the eastern part of the simulation area [25]. TRANSIMS runs (typical and non-typical) generate also a higher traffic volume on Alpha Rd than PAMINA (Fig. 2 (c)) and the Queueing model (Fig. 2 (d)) while the main difference between the last two simulations consists of higher traffic on the LBJ west-bound frontage road for PAMINA. As explained previously, the observed data, which serves as reference in each figure, shows more traffic in the north-south direction than the simulated data.

8 Comparisons using turn counts

Figs. 3 (a)–(d) compare *turn count* observations to simulation results. This means that approaching flows are now differentiated by turning directions. Since this type of information was not available for the NCTCOG assignment, we show results only for the TRANSIMS simulations.

Due to smaller aggregation, the turn counts show larger differences between the three simulations and the observed data than the approach counts. Nevertheless, the simulation results again look remarkably similar to each other. In general, the simulations underestimate flow in north-south direction, which can be traced back to the 1990 demand matrix as explained above.

Of the flow that is left, we have a tendency to distribute too much of it to minor streets. As a result, we also have, as a tendency, too many turns into minor streets and not enough into major streets. Assignment of demand to routes is done in the the routing module, though, and not in the micro-simulations.

Last, there seem to be more differences between MS1 (TRANSIMS) on one hand, and MS2 (PAMINA) and MS3 (QM) on the other hand. A possible explanation for this is that MS1 (TRANSIMS) has an explicit lane changing logic for plan following. Vehicles that want to make a turn all get into the same lane, *not* blocking traffic for other directions. The same is not true for MS2 and MS3.

A related issue is that, according to the last paragraph, MS1 generates different link travel times for different turning directions, but the planner module that was used does not use this information. That is, for re-planning we aggregate link travel times over *all* vehicles that leave a link, independent of the turning directions. It should be a question of future research how much difference it would make to employ a routing that includes such movement-dependent penalities. From our own experience (not shown), larger differences between the different simulations should be expected.

9 Discussion

The data consistency problem caused by the extension of the Dallas North Tollway is not an unlucky coincidence, but it is a generic problem in the field: The geographical areas that are most interesting for studies are the congested ones, and these are also the areas that change fastest. This clearly points to the need of a methodology that can deal with such changes in a systematic way. Thus, this becomes an example of why it is necessary to go beyond trip tables and generate demand from demographics via activities [5, 3]. Once that demand generation process is sufficiently understood, so the hope is, the additional trips caused by adding a road would be generated by the model so that the model would automatically react correctly and consistently to such infrastructure changes.

Given this data inconsistency problem, it is unclear how much could be learned by further studies. For the simulations, we could certainly change the network back to pre-extension status. However, that would only allow consistent comparison with the NCTCOG study, and this comparison for itself is meaningless. 1996 trip tables, which would allow a systematic comparison to the turn counts, were not available to us. In addition, the restriction of the simulation to the study area causes all kinds of boundary effects. For example, in reality access to the Tollway is really slow north of the simulated area; yet, since the re-planner does not pick this up because it is outside the simulated area, we probably route too much traffic via the Tollway.

Yet, a lot can be learnt from this study. First, judging from the admittedly limited data material, using iterated micro-simulations does not generate obviously more unrealistic flow results than traditional assignment. This in itself is good news since it indicates that microsimulation-based assignment will already in the near future be useful for policy decisions, plus it should improve further over time. Second, from the fact that the results of all three micro-simulations look similar, two conclusions can be drawn: (i) Deviations from reality do currently *not* come so much from the micro-simulation technology, but from the demand generation and from the route assignment part. (ii) Since the simplest micro-simulation technology that we employed (the QM model) generates similar results than the more extensive technologies, one can make a case that for the near-term future simplified micro-simulations may be good enough to drive research and maybe even for lower-fidelity studies. As mentioned above, this could change once one reports the link travel times dependent on the turning directions.

10 Summary

TRANSIMS assigns traffic to links via multiple iterations between micro-simulation and re-planner. The (re-)planner generates deterministic routes for each traveller; the micro-simulation executes these routes and remembers time-dependent link delays; the (re-)planner then generates new routes for a fraction of the population (using fastest path); etc. We showed results from a study where we used the same origin-destination-matrix and the same (re-)planner but three different micro-simulations. The micro-simulations were different in the level of reality that they contained, for example in the treatment of signal phasing, turn pockets, and lane connectivity for turning movements.

We used field observations of turn counts at intersections as a baseline for comparisons. A major problem with these field observations was that they were inconsistent with the data material that was driving the simulations. However, we argued that this will often be the case and rather than deplore the state of the data it is necessary to continue and increase work on technologies that can handle these restrictions.

The results of the micro-simulations turned out to be remarkably similar, indicating that (i) deviations from reality are currently more caused by the demand generation and by the routing than by the micro-simulations, and (ii) that, given the current knowledge and technology in demand generation and in routing, very simple micro-simulations may be sufficient for many questions in the near future.

For comparison, we showed the results of a traditional assignment for the same problem. The results indicate that microsimulation-based assignment, even given current technology, is not inferior to traditional assignment.

References

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

[2] M. Van Aerde, M. Baker, and J. Stewart. Weaving capacity sensitivity analysis using the INTEGRATION model. preprint.

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

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

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

[6] 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.

[7] R.J. Beckman et al. TRANSIMS–Release 1.0 – The Dallas-Fort Worth case study. Technical Report 97-4502, 1997.

[8] 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.

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

[10] 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.

[11] H.J. Herrmann, E. Flekkoy, K. Nagel, G. Peng, and G. Ristov. Density waves in granular flow, 1994.

[12] M. Rickert. *Traffic simulation on distributed memory computers.* PhD thesis, University of Cologne, Cologne, Germany, 1998.

[13] K. Nagel. Freeway traffic, cellular automata, and some (self-organizing) criticality. In R.A. de Groot and J. Nadrchal, editors, *Physics Computing '92*, page 419, Prague, 1992. World Scientific.

[14] K. Nagel and M. Schreckenberg. A cellular automaton model for freeway traffic. *Journal de Physique I France*, 2:2221, 1992.

[15] K. Nagel, D.E. Wolf, P. Wagner, and P.M. Simon. Two-lane traffic rules for cellular automata: A systematic approach. *Physical Review E*, 58(2):1425–1437, 1998.

[16] C. Gawron. An iterative algorithm to determine the dynamic user equilibrium in a traffic simulation model. pages 469–474.

[17] 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.

[18] H.P. Simão and W.B. Powell. Numerical methods for simulating transient, stochastic queueing networks. *Transportation Science*, 26:296, 1992.

[19] C.L. Barrett et al. TRANSIMS user notebook version 1.0. Technical Report 98-848, 1998.

[20] M. Rickert and P. Wagner. Parallel real-time implementation of large-scale, route-plan-driven traffic simulation. *International Journal of Modern Physics C*, 7(2):133–153, 1996.

[21] M. Rickert, P. Wagner, and Ch. Gawron. Real-time traffic simulation of the German Autobahn Network. In F. Hoßfeld, E. Maehle, and E.W. Mayer, editors, *Proceedings of the 4th PASA Workshop*. World Scientific, Singapore, 1996.

[22] Alfredo Garcia and Robert L. Smith. On the convergence of iterative routing-assignment procedures in dynamic traffic networks. Technical Report 95-26, University of Michigan Department of Industrial and Operations Engineering, Ann Arbor MI 48109, November 1995.

[23] David E. Kaufman, Robert L. Smith, and Karl E. Wunderlich. User-equilibrium properties of fixed points in iterative dynamic routing/assignment methods. *Transportation Research C*, forthcoming.

[24] J. Bottom, M. Ben-Akiva, M. Bierlaire, and I. Chabini. Generation of consistent anticipatory route guidance.

[25] K. Nagel. Experiences with iterated traffic microsimulations in Dallas. pages 199–214.

[26] K. Nagel, M. Rickert, and P.M. Simon. The dynamics of iterated transportation simulations. Technical Report 98-2168, 1998. Presented at the 3rd TRIannual Symposium on Transportation ANalysis (TRISTAN-III) in San Juan, Puerto Rico.

[27] Major investment study for the I.H. 635 (LBJ freeway) corridor. Texas Department of Transportation, Dallas, Texas. See `www.dot.state.tx.us`, 1996.

Figure 1: Comparison between the NCTCOG 1990 assignment and the 1996 approach counts. Bars at intersections give the hourly approach volumes at intersections according to the scale shown in the figure. The wider, gray bars denote the flows generated by the model. Black narrower bars denote observations. When a black bar is "sticking out", it means that the model underestimated flows by that amount. When a black bar is shorter than the gray bar, it means that the model overestimated flows by that amount. – For four intersections in the middle of the figure no assignment values were available so that only the observations are shown.

Figure 2: (a) Comparison between a "typical" run of MS1 (TRANSIMS) and the 1996 approach counts. (b) Comparison between a "non-typical" run of MS1 (TRANSIMS) and the 1996 approach counts. (c) Comparison between MS2 (PAMINA) and the 1996 approach counts. (d) Comparison between MS3 (QM) and the 1996 approach counts.

Figure 3: Comparisons using turn counts, that is, approaching cars are differentiated by turning direction. In consequence, instead of one major bar we now have three major bars per approach direction, one for left turns, one for goint straight, and one for right turns. (a) Comparison between a "typical" run of MS1 (TRANSIMS) and 1996 turn counts. (b) Comparison between a "non-typical" run of MS1 and 1996 turn counts. (c) Comparison between MS2 (PAMINA) and 1996 turn counts. (d) Comparison between MS3 (QM) and 1996 turn counts.


Alpha
Peterson
Tollway
Frontage Road
James Temple
Valley View Mall
2500 Cars
0 Cars
Preston
Galleria
Frontage Road
LBJ Freeway
obs_1h.gui
nctcog_model-1h.gui
8:00 <= t < 9:00

Fig. 1


Alpha
Peterson
Tollway
Frontage Road
James Temple
Valley View Mall
2500 Cars
0 Cars
Preston
Galleria
Frontage Road
LBJ Freeway
obs_1h.gui
microsim.14-plan-seed-456-1h.gui
8:00 <= t < 9:00

Fig. 2 (a)

Fig. 2 (b)


Alpha
Peterson
Tollway
Frontage Road
James Temple
Valley View Mall
2500 Cars
0 Cars
Preston
Galleria
Frontage Road
LBJ Freeway
obs_1h.gui
microsim.14-plan-seed-123-1h.gui
8:00 <= t < 9:00

Alpha
Peterson
Tollway
Frontage Road
James Temple
Valley View Mall
Preston
2500 Cars
0 Cars
Galleria
Frontage Road
LBJ Freeway
obs_1h.gui
mr-run11-it60-1h.gui
8:00 <= t < 9:00

Fig. 2 (c)

Fig. 2 (d)

Alpha
Peterson
Tollway
Frontage Road
James Temple
Valley View Mall
Preston
2500 Cars
0 Cars
Galleria
Frontage Road
LBJ Freeway
obs_1h.gui
queue-1h.gui
8:00 <= t < 9:00

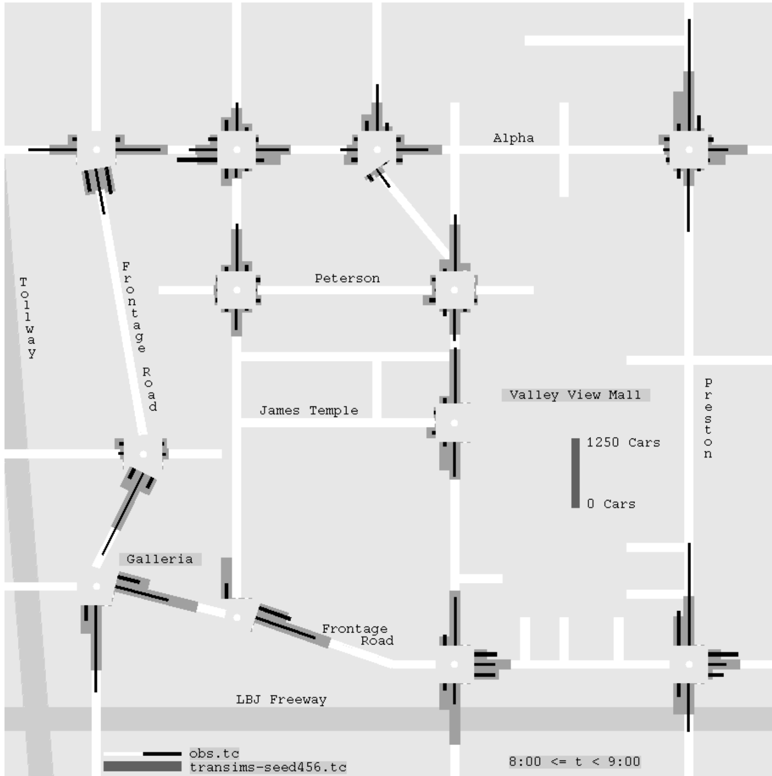
Alpha
Peterson
Tollway
Frontage Road
James Temple
Valley View Mall
Preston
1250 Cars
0 Cars
Galleria
Frontage Road
LBJ Freeway
obs.tc
transims-seed456.tc
8:00 <= t < 9:00

Fig. 3 (a)

Fig. 3 (b)

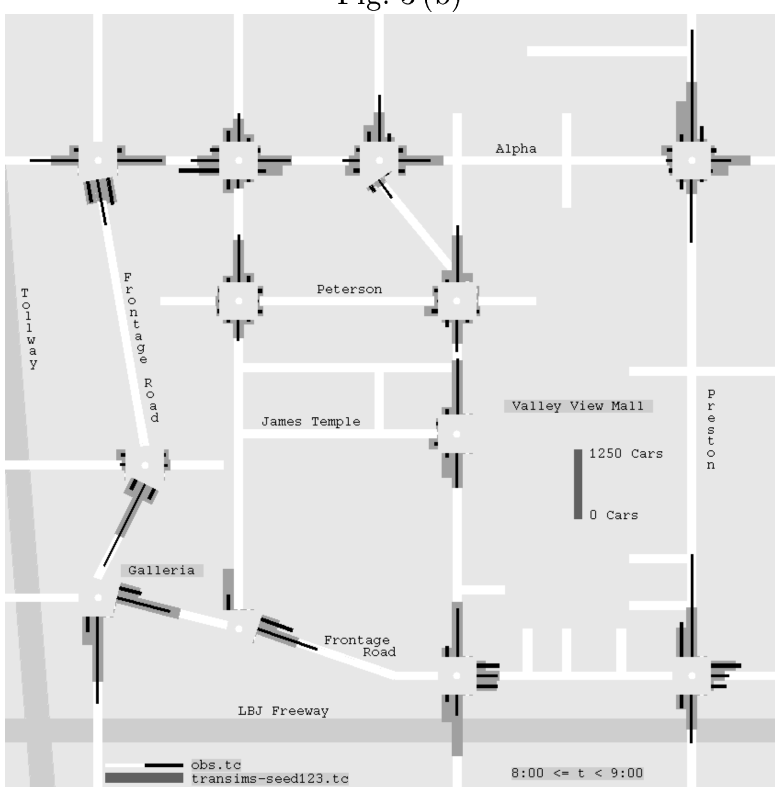
Alpha
Peterson
Tollway
Frontage Road
James Temple
Valley View Mall
Preston
1250 Cars
0 Cars
Galleria
Frontage Road
LBJ Freeway
obs.tc
transims-seed123.tc
8:00 <= t < 9:00

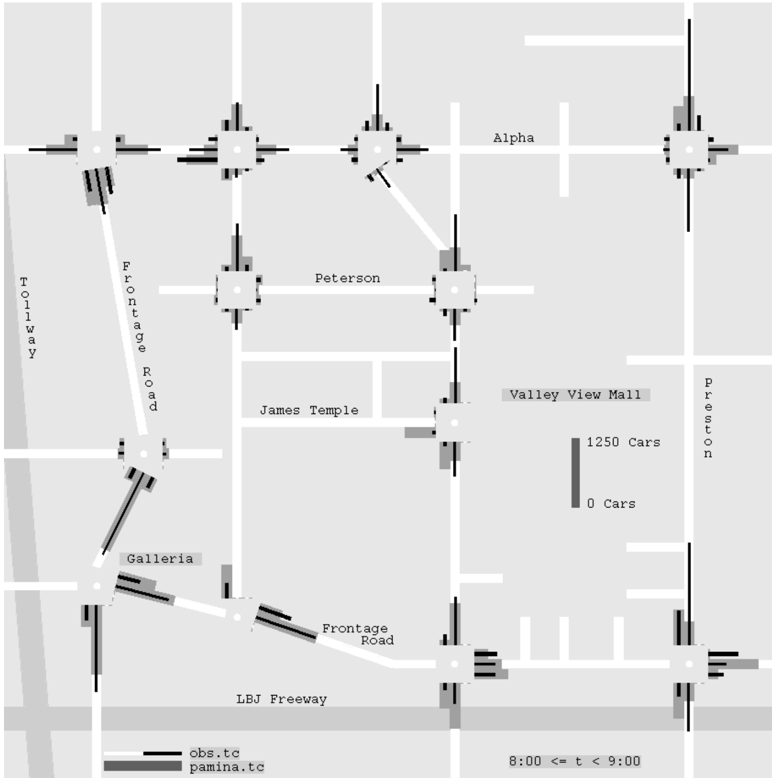
Alpha
Peterson
Tollway
Frontage Road
James Temple
Valley View Mall
Preston
1250 Cars
0 Cars
Galleria
Frontage Road
LBJ Freeway
obs.tc
pamina.tc
8:00 <= t < 9:00

Fig. 3 (c)

Fig. 3 (d)

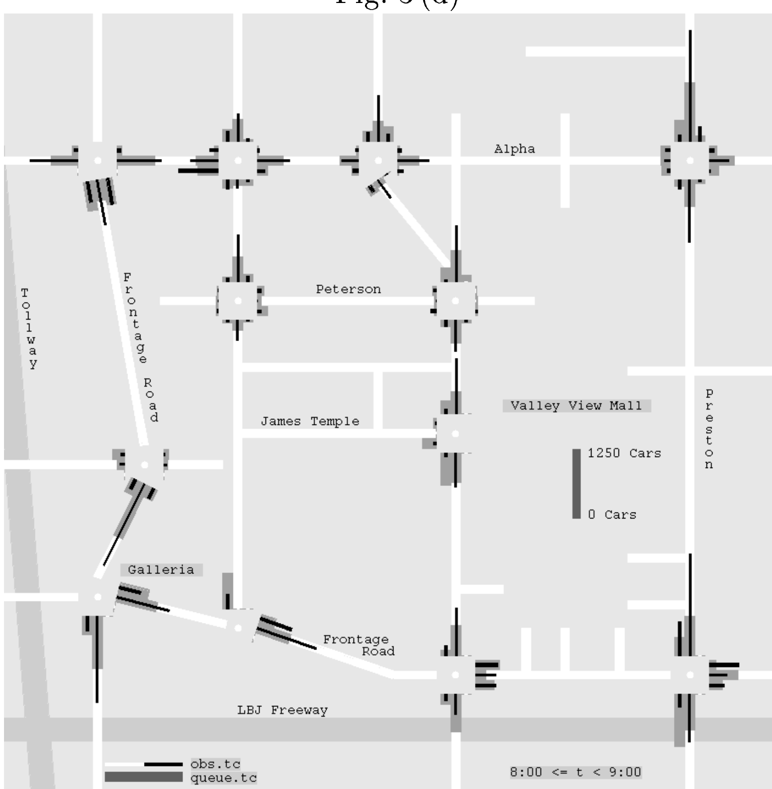
Alpha
Peterson
Tollway
Frontage Road
James Temple
Valley View Mall
Preston
1250 Cars
0 Cars
Galleria
Frontage Road
LBJ Freeway
obs.tc
queue.tc
8:00 <= t < 9:00