Online Traffic State Estimation with Multi-Agent Simulations

Gunnar Flötteröd, Kai Nagel
Berlin Institute of Technology
Sek SG 12, Salzufer 17-19
10587 Berlin, Germany
{floetteroed,nagel}@vsp.tu-berlin.de

Abstract: This article describes a novel approach to online traffic state estimation with multi-agent simulations. The presented method identifies individual-level motorist behavior from aggregate measurements of flows, densities or velocities that are obtained at a limited set of network locations. Behavioral dimensions that can be estimated range from single route choice to the selection of full-day plans. The estimation logic only requires a simulation-based representation of the behavioral dimensions under consideration in order to simulate the most plausible driver behavior given the additional information contained in the sensor data. No mathematical description of the behavioral driver model is needed such that the broad modeling capabilities of a multi-agent simulation can be fully exploited.

Key Words: traffic state estimation, online traffic monitoring, disaggregate demand calibration, traffic counts, multi-agent simulation, microsimulation

1 Introduction

There are basically two aspects of a traffic system that may be subject to state estimation: the demand model and the supply model. Traditionally, the demand is represented in terms of an OD (origin/destination) matrix, which is loaded on the network by the supply model. Our work focuses on the online calibration of the demand model in a fully disaggregate manner. Of course, there also are uncertainties in every supply model, which need to be accounted for in a comprehensive traffic monitoring system.

The estimation of time-dependent OD matrices can be seen as the aggregate counter piece to our disaggregate demand calibration approach. OD matrix calibration from traffic counts has received much attention in the literature [4, 9, 19, 29, 30]. The nonlinear and dynamical nature of traffic has successfully been accounted for [10, 11, 20, 32, 33], and several online implementations have been reported [1, 2, 31, 34].

An OD matrix estimator captures various behavioral aspects on an aggregate level: Since a time-dependent OD matrix maps (origin, destination, departure time) tuples on demand levels, it directly represents destination and departure time choice. A motorist OD matrix reflects mode choice at least in terms of decisions for or against the vehicular mode. Route choice, however, constitutes no additional degree of freedom but is a function of the demand that is defined by the traffic assignment procedure. Path flow estimators constitute a notable exception to this, yet only in a macroscopic and mostly static setting [5, 6, 7, 22, 23, 27, 28].

The calibration of disaggregate supply models from aggregate sensor data has in fact been widely addressed in the literature, e.g. [12, 13, 14, 17, 18, 24, 25]. However, to the best of our knowledge, no such effort has yet been made to calibrate the demand model in a likewise disaggregate manner.

2 Methodology

2.1 Multi-agent simulations

The model deployed by our traffic state estimation system is a multi-agent simulation. This is not to say that our methodology is constrained to this particular simulation technique. We rather consider the multi-agent approach to be the most challenging one from an estimation point of view because of its extremely broad modeling capabilities, and thus expect our methodology to be applicable also in conjunction with more traditional model types.

The multi-agent approach is characterized by the fully disaggregate representation of travelers throughout the entire modeling process. This approach is attractive in the traffic domain since it appears natural to represent every traveler by a software object, to put these individual models into a representation of the physical world of mobility, and to observe the resulting mobility patterns.

Multi-agent simulation can go beyond other simulation methods by including travelers' goals and commitments into the modeling. For example, it is possible to differentiate between a delayed person with a free evening and a delayed person with a time-restricted day-care pick-up. A multi-agent simulation for transportation typically consists of the following modules [3, 21, 26]:

- A synthetic population generation module generates, from demographic data, a synthetic population that, in all its statistical aspects, corresponds to the real population under investigation, while at the same time preserving privacy.
- An activity-based demand generation module generates, for each member of the synthetic population, complete daily plans including a sequence of activities (such as home, work, shop, leisure), activity locations, and a temporal schedule. Consecutive activities at different locations generate the demand for travel.
- A router module computes how that demand is actually executed on the network, possibly including mode choice. At this point, all synthetic travelers have plans that describe what they intend to do.
- Another module puts the synthetic travelers on a simulated version of the physical network and has them execute their plans simultaneously. The physical interaction in that system generates congestion. Depending on the specific focus, this module has different names: supply simulation, network loading, traffic flow simulation.

It is not possible to compute the system in the linear way indicated above since plans depend on congestion but congestion is a consequence of the plans. This is solved by iterations that can be seen as modeling human day-to-day learning, and which eventually lead to some kind of an equilibrium situation.

In a telematics setting where drivers act based on incomplete information and, often enough, not as one would rationally expect, an equilibrium-based modeling approach is inappropriate. The multi-agent technique can also handle such a setting, basically by omitting the iterations: All agents start their simulated day with some background knowledge about the expected traffic conditions. This knowledge has been generated beforehand, e.g. by an iterative planning simulation. If an unforeseeable event occurs within the simulated day, the agents react to this event only based on their current level of information, as one would expect in reality.

2.2 Estimation logic

The complexity of a multi-agent simulation renders the design of a flexibly applicable estimator a nontrivial task. In order to be compatible with the proposed estimator, a traffic simulation system must be separable into the components shown in Figure 1. Most of the subsequently employed terminology is adopted from [8].

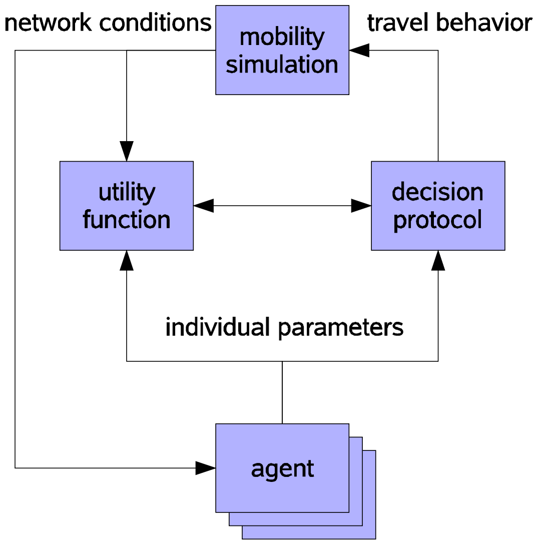

Figure 1: Simulation

The **mobility simulation** moves individual vehicles along their chosen routes through the road network. All physical interactions occur within this component. The trip sequence of every vehicle is chosen by an individual **agent** that represents the driver of that vehicle. The decision making of an agent is realized by two further components. Whenever a decision is required, the agent provides these components with its individual parameters.

- The **utility function** provides an individually parameterized map from network conditions on the systematic utility of any behavioral alternative available to the agent.
- The (likewise individually parameterized) **decision protocol** probabilistically generates a single decision based on this utility information.

Estimation is based on reasonable mathematical inference [15, 16] yet follows a simple technical logic. As illustrated in Figure 2, the simulation structure is not changed at all. An **estimator** component is inserted between the decision protocol and the remaining simulation system. It is implemented transparently in that it provides unmodified interfaces to both, the decision protocol and the remaining system. The estimator compares the output of the mobility simulation to measurements from a surveillance system. Based on this comparison, it alters the data and control flow around the decision protocol such that the resulting agent behavior is most plausible given the measurements.

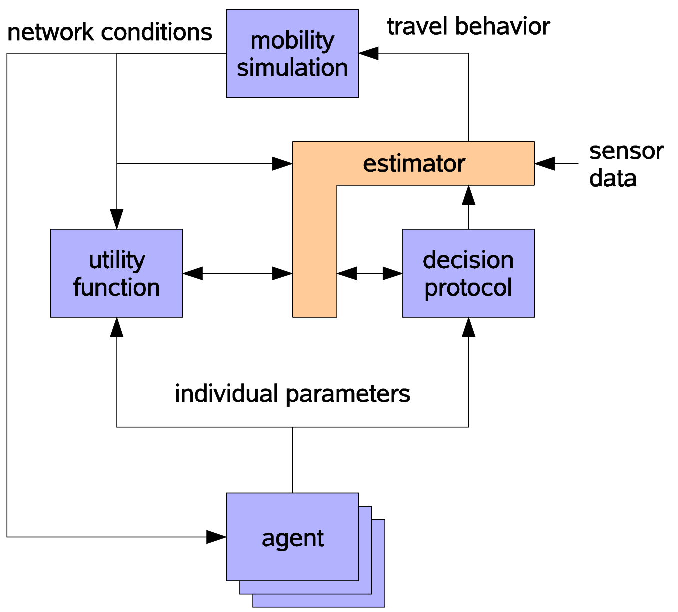

Figure 2: Estimation

Two small route choice examples illustrate how this minor system extension allows adjusting simulated behavior.

- If the surveillance system observes a traffic jam where there is none in the simulation, the estimator increases the systematic utility of the according links until the agents start to favor these links and create the congestion as observed in reality. Vice versa, if there is congestion in the simulation but not in reality, the estimator decreases the involved links' utility until agents start to avoid the critical area.
- Likewise, the estimator can encourage a certain behavioral pattern by asking the decision protocol to draw several alternatives in identical conditions for each agent. From this set of options the estimator then passes only those decisions on to the mobility simulation that are most plausible given the observed measurements.

Either approach accesses only a subset of the interfaces touched by the estimator in Figure 2. This further relaxes the structural requirements on the simulation system. The apparent simplicity of this approach is confronted with (i) the difficulties to relate aggregate measurements and individual behavior through nonlinear traffic flow dynamics on large networks of general topology, and (ii) the intention to be compatible with a broad variety of behavioral implementations.

3 Experimental results

3.1 Setup

The presented experiments are conducted on a 2459-link network that represents the major road network of Greater Berlin. The synthetic population is of size 206'353. The experiments are constrained to the time span from 6 to 9 am. This interval exhibits the most variable traffic conditions because of the morning rush hour.

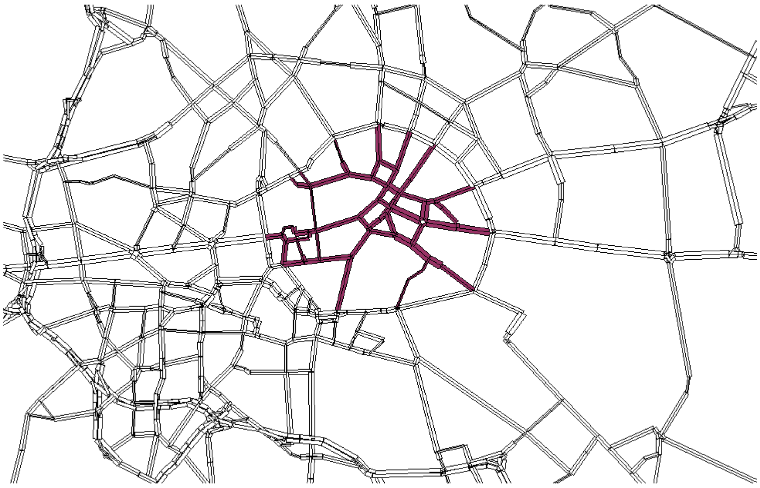

Figure 3: Toll zone

A hypothetical time-independent toll of 0.24 EUR/km is charged in the city center shown in Figure 3, and no toll is charged outside of this area. The sole behavioral degree of freedom considered here is route choice. The (dis)utility of a route is additively comprised of (i) the travel time on that route and (ii) the toll accumulated along that route. In order to combine these two utility components into a single number, the travel time is turned into a monetary quantity by multiplication with a **value of time** (VOT) parameter.

The first day after implementation of the toll is considered. In this situation, the drivers are aware of the typical travel times without toll and of the toll itself. However, they have not yet learned the alterations in traffic conditions that result from other travelers' changed behavior in response to the toll. In consequence, the presumably most advantageous route choice for many drivers that so far have traversed the toll area is now to avoid it but to bypass it as sharply as possible in order to minimize the resulting increase in travel time. This would cause an unforeseeable congestion on the roads that immediately encircle the toll zone. This behavior, generated from (i) network conditions from before the toll, (ii) the toll, and (iii) a 12 EUR/h VOT assumption, makes up the **prior scenario**, corresponding to the prediction that the estimation system would make in the absence of measurements.

Synthetic reality, in contrast, is assumed to be a scenario where the toll is effectively ignored. That is, the severe congestion around the toll zone that is predicted by simulation of the prior scenario does not occur in the synthetic reality.

Keeping in mind that only the first day after installation of the toll is simulated, such a behavior may either result from unawareness or from curiosity about the involved technical installations. It must be stressed that the purpose of these experiments is to sound the capabilities of the estimator, not to discuss road pricing issues themselves.

The task of the estimation system now is to reconstruct as much as possible of the synthetic reality, given limited measurements from that synthetic reality, and given the prior scenario. Flow measurements, i.e. traffic counts at road cross-sections per time interval, are used in all experiments as synthetically generated sensor data. This data is generated from 50 flow sensors in the synthetic reality. All such data is averaged in 5 minute time bins.

Network-wide occupancy information is used for validation purposes. That is, global network conditions are compared in terms of the average number of vehicles on every link in every 5-minute time bin from 6 to 9 am. This comparison is quantified in terms of a root mean square error measure, denoted as the **RMS error**, which basically represents the Euclidean distance between a simulated/estimated occupancy vector and that of the synthetic reality.

Experiments are conducted in offline and simulated online conditions. In offline conditions, a set of beforehand collected measurements is processed "en block". In a telematics context, this is useful for the ex post analysis of a particular day. In contrast, the online estimator runs in a rolling horizon mode where the estimation of the traffic state for a certain point in time has only measurements from earlier times available. This setting is characteristic for a continuous traffic monitoring problem.

A more detailed description of these experiments can be found in [16].

3.2 Offline experiments

To begin with, the rolling horizon mode is not facilitated and a sequence of offline estimations is run over the entire 6 to 9 am time period. These experiments focus on the quality of calibration. The next section will present experiments in simulated online conditions, investigate the estimator's real time capabilities, and conclude about the scenario size its current implementation can handle.

Figure 4 shows the resulting RMS validation errors. They are drawn over a range of w_prior parameter values. w_prior represents the analyst's belief in the correctness of the simulated prior network conditions. The smaller w_prior, the lower this belief, and the more weight is put on the reproduction of the sensor data. Vice versa, a very large w_prior effectively ignores the sensor data and only reproduces the prior scenario.

The blue dots represent the RMS errors between three estimation runs per w_prior value and the synthetic reality. The RMS errors between three plain simulations of the 12 EUR/h prior scenario and the synthetic reality are drawn in red. Although a plain simulation of the prior scenario corresponds to an infinite w_prior, the simulation results are re-drawn over each w_prior value for ease of comparison. Since the simulation and estimation results are very stable over repeated runs, most dots per w_prior value are located on top of each other and cannot be visually distinguished.

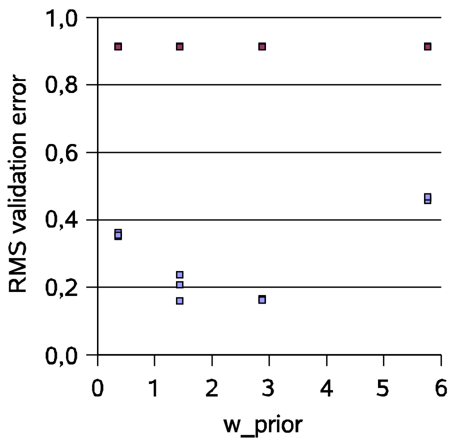

Figure 4: Offline estimation results

There is a non-monotonous relation between w_prior and the RMS error. As w_prior grows, the measurements' influence vanishes and the estimation quality gracefully deteriorates towards that of plain simulation of the prior scenario. However, as w_prior decreases, a minimum RMS error value is encountered after which a further decrease of w_prior results in an increased RMS error. A plausible explanation for this increase is overfitting to the 50 measurement locations, reducing the estimation quality where no measurements are available. The attained minimum RMS error reflects the estimator's ability to spatiotemporally extrapolate the available flow measurements. For w_prior=2.88, the estimator improves the RMS error by 82% over a plain simulation of the prior scenario. The severe congestion of the 12 EUR/h VOT prior scenario that does not occur in the synthetic reality is successfully prevented by the estimator.

3.3 Online experiments

A rolling horizon logic is implemented that runs the estimator in simulated online conditions. The time period of investigation still is 6 to 9 am. However, while one iteration of an offline estimator facilitates all measurements from this interval at once, online conditions imply that the measurements become available bit by bit as the simulated real time proceeds.

Online estimation starts at 6:30 simulated real time. Only measurements until this moment are available. During this first **estimation period**, only a simulation from 6:00 to 6:30 is iteratively adjusted. After a prespecified number of iterations, the simulated real time is advanced to 6:35, the most recent simulation is continued until 7:00 to evaluate the estimator's predictive capabilities, measurements from 6:30 to 6:35 become available, and the next estimation period from 6:05 to 6:35 begins. All driver behavior until 6:05 is now fixed according to the most recent iteration of the previous estimation period.

Note that the iterative nature of the estimator has not been considered in the previous offline experiments, although those experiments used iterations as well. In online conditions, however, the number of computationally intensive iterations that is needed until a result is generated becomes a crucial factor for the estimator's applicability.

A prior weight of w_prior=2.88 is maintained in all runs since this setting achieved the best results in the offline experiments. Figure 5 provides separate results for every 30-minute estimation period ending at 7 through 9 am. The blue bars represent (from left to right) RMS errors obtained at the end of 5, 10, 20, 30, 40, and 50 iterations per estimation period. They are drawn on top of red RMS error bars that result from plain rolling horizon simulations with respective iteration numbers. These simulations follow an identical rolling horizon logic as the estimator, only that the measurements are not accounted for, i.e. effectively the estimator predicts the prior scenario.

Both, the estimation errors and the errors of the prior rise over time as the traffic volumes increase in the morning rush hour. A pronounced difference between the prior and estimation can be observed as the congestion around the toll zone becomes severe in the prior scenario. Overall, the estimator reduces the RMS error by up to 70% in the later periods. Conducting only 5 or 10 iterations per estimation period result in lower improvements when compared to 20 iterations and more. However, running beyond 20 iterations yields only marginal improvements.

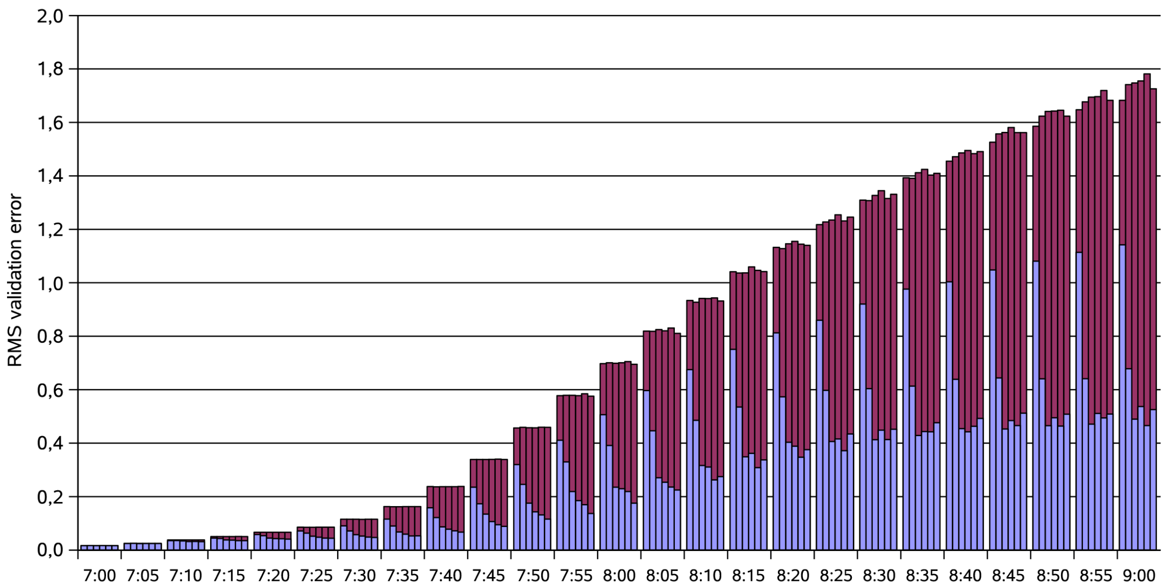

Figure 5: Online estimation results

Figure 6 shows the same setup of RMS errors as before, only that now the average *prediction* errors over a 0 to 30 minute time interval are given. This and the previous diagram match temporally in the following way. An estimation error drawn e.g. over the 8:30 label was generated at this particular time and thus applies to the interval from 8:00 to 8:30. A prediction error that is drawn over the 8:30 label was generated at 8:00 for a 30 minute prediction window and consequently applies to the same interval. A comparison of both figures yields the expected diagnosis that the estimation quality is generally higher than the prediction quality. However, the estimation-based prediction is clearly better than the plain simulation. Again, the prediction results for 5 and 10 iterations per estimation period are inferior when compared to those with 20 iterations and more. The computational effort of executing more than 20 iterations per estimation period does not result in significantly improved predictions. Overall, the estimator reduces the RMS prediction error by 50% to 60% in the later time periods.

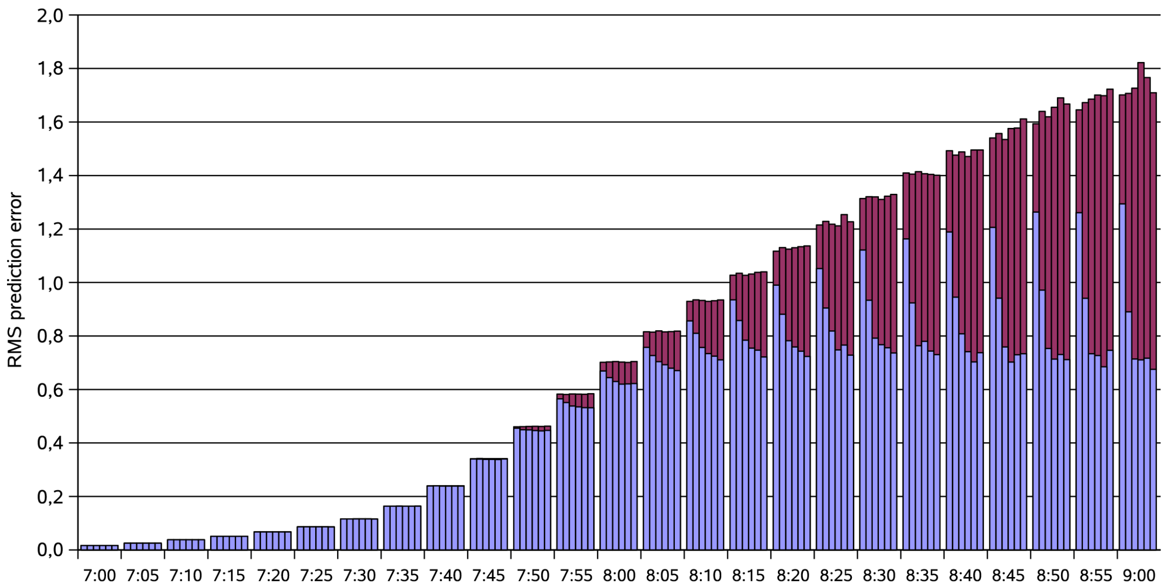

Figure 6: Online estimation results (prediction)

The current implementation of the estimator accomplishes 6 iterations per 5-minute interval in the given scenario. That is, near-optimal results require another estimation speedup of 3 to 4. Given the considered problem's size, this is an encouraging result, and even with only 6 iterations per 5 minutes, a suboptimal estimation still yields substantial improvements when compared to the prior scenario. Since the computational effort rises at least linearly with the network and population size, a 600+ link scenario with 50'000+ agents is immediately approachable by the current implementation in real time.

4 Acknowledgements

This project was funded in part by the German research society DFG under the grant “State estimation for traffic simulations as coarse grained systems”. Significant amounts of computing time on the computing cluster of TU Berlin's mathematical faculty are gratefully acknowledged.

5 Summary and outlook

We presented a novel approach to the online calibration of a microscopic demand simulator from aggregate sensor data. The conceptual background of our method was briefly outlined, and experimental results were presented which indicate the practical relevance of our work.

Our ongoing research focuses on the further relaxation of the modeling assumptions we currently make when applying the estimator. So far, we have successfully decoupled the estimation methodology from the particular nature of the behavioral simulation component. An important aspect of our future work is the likewise flexible treatment of a general simulation-based and stochastic mobility simulation.

References

[1] C. Antoniou. On-line Calibration for Dynamic Traffic Assignment. PhD thesis, Massachusetts Institute of Technology, 2004.

[2] K. Ashok. Estimation and Prediction of Time-Dependent Origin-Destination Flows. PhD thesis, Massachusetts Institute of Technology, 1996.

[3] M. Balmer, B. Raney, and K. Nagel. Adjustment of activity timing and duration in an agent-based traffic flow simulation. In H.J.P. Timmermans, editor, Progress in activity-based analysis, pages 91-114. Elsevier, Oxford, UK, 2005.

[4] M.G.H. Bell. The estimation of origin-destination matrices by constrained generalised least squares. Transportation Research Part B, 25:13-22, 1991.

[5] M.G.H. Bell and S. Grosso. Estimating path flows from traffic counts. In W. Brilon, F. Huber, M. Schreckenberg, and H. Wallentowitz, editors, Traffic and mobility: simulation-economics-environment, pages 85-102. Springer, 1999.

[6] M.G.H. Bell, W.H.K. Lam, and Y. Iida. A time-dependent multi-class path flow estimator. In J.-B. Lesort, editor, Proceedings of the 13th International Symposium on Transportation and Traffic Theory, Lyon, France, July 1996. Pergamon.

[7] M.G.H. Bell, C.M. Shield, F. Busch, and G. Kruse. A stochastic user equilibrium path flow estimator. Transportation Research Part C, 5(3/4):197-210, 1997.

[8] J.L. Bowman and M. Ben-Akiva. Activity-based travel forecasting. In Activity-Based Travel Forecasting Conference: Summary, Recommendations and Compendium of Papers, New Orleans, Louisiana, June 1996. See http://tmip.fhwa.dot.gov/clearinghouse/.

[9] E. Cascetta. Estimation of trip matrices from traffic counts and survey data: a generalised least squares estimator. Transportation Research Part B, 18:289-299, 1984.

[10] E. Cascetta, D. Inaudi, and G. Marquis. Dynamic estimators of origin-destination matrices using traffic counts. Transportation Science, 27:363-373, 1993.

[11] E. Cascetta and N.N. Posterino. Fixed point approaches to the estimation of o/d matrices using traffic counts on congested networks. Transportation Science, 35(2):134-147, 2001.

[12] L. Chu, H.X. Liu, J.-S. Oh, and W. Recker. A calibration procedure for microscopic traffic simulation. In Proceedings of the 6th IEEE Intelligent Transportation Systems Conference, Shanghai, October 3-6 2003.

[13] S.L. Cohen. An approach to calibration and validation of traffic simulation models. In Proceedings of the 83. Annual Meeting of the Transportation Research Board, Washington DC, USA, 2004.

[14] R. Dowling, A. Skarbadonis, J. Halkias, G.M. Hale, and G. Zammit. Guidlines for calibration of microsimulation models: Framework and applications. In Proceedings of the 83. Annual Meeting of the Transportation Research Board, Washington DC, USA, 2004.

[15] G. Flötteröd. A fully disaggregate approach to the calibration of simulation-based DTA. In DTA2008 International Symposium on Dynamic Traffic Assignment, Leuven, Belgium, June 2008.

[16] G. Flötteröd. Traffic State Estimation with Multi-Agent Simulations. PhD thesis, Berlin Institute of Technology, Germany, 2008.
[17] S.-J. Kim, W. Kim, and L.R. Rilett. Calibration of micro-simulation models using non-parametric statistical techniques. In Proceedings of the 84. Annual Meeting of the Transportation Research Board, Washington DC, USA, 2005.
[18] D.-H. Lee, X. Yang, and P. Chandrasekar. Parameter calibration for PARAMICS using genetic algorithm. In Proceedings of the 80. Annual Meeting of the Transportation Research Board, Washington DC, USA, January 2001.
[19] M. Maher. Inferences on trip matrices from observations on link volumes: a Bayesian statistical approach. Transportation Research Part B, 17:435-447, 1983.
[20] M.J. Maher, X. Zhang, and D. Van Vliet. A bi-level programming approach for trip matrix estimation and traffic control problems with stochastic user equilibrium link flows. Transportation Research Part B, 35:23-40, 2001.
[21] K. Nagel and F. Marchal. Computational methods for multi-agent simulations of travel behaviour. In Kay W. Axhausen, editor, Moving through nets: The physical and social dimensions of travel. Selected Papers from the 10th international conference on travel behaviour research. Elsevier, 2007.
[22] Y. Nie and D.-H. Lee. An uncoupled method for the equilibrium-based linear path flow estimator for origin-destination trip matrices. Transportation Research Record, 1783:72-79, 2002.
[23] Y. Nie, H.M. Zhang, and W.W. Recker. Inferring origin-destination trip matrices with a decoupled GLS path flow estimator. Transportation Research Part B, 39:497-518, 2005.
[24] B. Park and H. Qi. Development and evaluation of simulation model calibration procedure. In Proceedings of the 84. Annual Meeting of the Transportation Research Board, Washington DC, USA, 2005.
[25] B. Park and J.D. Schneeberger. Microscopic simulation model calibration and validation: A case study of VISSIM for a coordinated actuated signal system. Transportation Research Record, 1856:185-192, 2003.
[26] B. Raney and K. Nagel. An improved framework for large-scale multi-agent simulations of travel behavior. In P. Rietveld, B. Jourquin, and K. Westin, editors, Towards better performing European Transportation Systems. Routledge, Oxon, UK, 2006.
[27] H.D. Sherali, A.Narayan, and R. Sivanandan. Estimation of origin-destination trip-tables based on a partial set of traffic link volumes. Transportation Research Part B, 37:815-836, 2003.
[28] H.D. Sherali, R. Sivanandan, and A.G. Hobeika. A linear programming approach for synthesizing origin-destination trip tables from link traffic volumes. Transportation Research Part B, 28(3):213-233, 1994.
[29] H. Spiess. A maximum likelihood model for estimating origin-destination models. Transportation Research Part B, 21:395-412, 1987.
[30] H. van Zuylen and L. G. Willumsen. The most likely trip matrix estimated from traffic counts. Transportation Research Part B, 14:281-293, 1980.
[31] Y. Wen, R. Balakrishna, M. Ben-Akiva, and S. Smith. Online deployment of dynamic traffic assignment: architecture and run-time management. IEE Proc. Intell. Transp. Syst., 153(1):76-84, March 2006.
[32] H. Yang. Heuristic algorithms for the bilevel origin/destination matrix estimation problem. Transportation Research Part B, 29:231-242, 1995.
[33] H. Yang, T. Sasaki, and Y. Iida. Estimation of origin-destination matrices from link traffic counts on congested networks. Transportation Research Part B, 26(6):417-434, 1992.
[34] X. Zhou. Dynamic origin-destination demand estimation and prediction for off-line and on-line dynamic traffic assignment operation. PhD thesis, University of Maryland, College Park, 2004.