

State estimation for multi-agent simulations of traffic

Gunnar Flötteröd, Kai Nagel

Technical University of Berlin

Transport Systems Planning and Transport Telematics Group

Salzufer 17-19, Sek SG12, 10587 Berlin

floetteroed@vsp.tu-berlin.de, nagel@vsp.tu-berlin.de

Abstract

This article describes a novel method for microsimulation-based traffic state estimation, which adjusts individual travelers' route and activity location choice to anonymous measurements e.g. of flows or velocities. While a discussion of the algorithm's rather mathematical functioning is omitted, the approach is clarified by means of an illustrative example. A second example of realistic size underlines the method's real world applicability and its real time capabilities.

1 Introduction

The problem of traffic monitoring and prediction has been considered by many researchers. Various approaches are data-driven (Huisken and Berkum, 2003; Kamarianakis and Prastacos, 2003; Zhou and Nelson, 2002), while others adjust structural models to real world measurements. The latter group can further be classified with respect to what quantities are estimated: Some consider the problem of estimating physical traffic flow properties such as densities, velocities, or flow parameters (Lorkowski and Wagner, 2005; Wang and Papageorgiou, 2005), while others (including this work) concentrate on the underlying demand itself and consider the physics of traffic flow as a dependent effect (Ashok, 1996). The second point of view is closer to the real problem's

structure, since traffic demand is the cause of road usage. Still, estimation of traffic demand and of network link related quantities are two aspects of the same problem, and ultimately should not be separated (Antoniou, 2004).

This article describes a method for traffic state estimation with multi-agent simulations. We combine a flexible but little formalized representation of individual mobility behavior as implemented in the MATSim project (MATSim [www](http://www.matsim.org) page) with well understood methods of system engineering (e.g. Kumar and Varaiya, 1986). This allows us to consider the problem of estimating agents' route and activity location choice in a Bayesian setting by combining for every agent an *a priori* activity plan for a given day with anonymous traffic *measurements* such as flows or densities obtained during this day into a most likely *a posteriori* plan.

Our work appears to be the first in this field which estimates fully individualized behavior from anonymous traffic measurements. The choice of this objective is justified by the observation that traffic demand results from heterogeneous individual mobility needs. Thus, no validated individualized knowledge should be aggregated away during the formalizing steps of setting up a mathematical estimation problem.

The remainder of this article is organized as follows. Section 2 is devoted to modeling and simulation. While the focus of this article is on behavioral issues, an introduction to the employed mobility simulation is given as well. In section 3, our technical approach of extending a pure simulation system with state estimation capabilities is explained. Section 4 describes the methodological aspects of our approach in terms of a synthetic example, while in section 5 preliminary results from a case study of realistic size are presented. Section 6 concludes the article and gives an outlook on future work.

2 Deterministic modeling and simulation

The traffic model consists of two interacting major components: A mobility simulation that describes the dynamics of traffic flow, and a behavioral model which represents spontaneous driver behavior in terms of route and activity location choice.

While both model components comply with the formal requirements of the estimation procedure described in section 3, this section is confined to a conceptual outline from a practical point of view.

2.1 Physical model of traffic flow

The physical model combines microscopic and macroscopic aspects. The representation of traffic flow dynamics is a fully macroscopic 1st order traffic flow model which runs in discrete time and space. The model permits linearization, which allows predicting the effect of small parameter variations without repeated simulations (Flötteröd and Nagel, 2005). In this way, it allows to systematically search for improved parameter sets given a certain objective. In the estimation application, this objective will be „better explanatory power for given measurements“. On the other hand, we also require the model to work on a microscopic level in order to allow for arbitrary behavioral heterogeneity in the driver population, which is difficult to deal with in a macroscopic way.

The fully macroscopic traffic flow model moves continuous flows according to macroscopic fundamental diagrams. However, at intersections the flow is split according to turning fractions which result from individual behavior: Whenever an individual vehicle starts a trip, it is put into the network and an equivalent amount of macroscopic flow is dismissed into the system. The vehicle then is moved across its current link according to the velocity field as it is defined by the macroscopic model. Since these velocities depend on the link's macroscopic occupancy, the vehicle's entrance effectively influences the

macroscopic traffic situation. At the link's downstream intersection, the vehicle is free to choose its next link according to its internal behavioral model. In order to synchronize the macroscopic flow with the individual behavior, these microscopic turning movements are counted, filtered, and normalized into macroscopic turning fractions. When the vehicle leaves the system, an appropriate amount of flow is removed as well at the exit point. Overall, the approach is similar to what is termed “smoothed particle hydrodynamics (SPH)” in physics (Gingold and Monaghan, 1977) or “mesoscopic modeling” in transport science (Ben-Akiva et al., 1998; DYNAMIT www page; Chang et al., 1985; DYNASMART www page; Schwerdtfeger, 1987), the main difference being that the model described here was designed with the explicit intention to obtain first derivatives from the model.

The interplay between both simulation aspects can thus be stated as: Massless microscopic vehicles float through the network according to fully macroscopic laws, while the intersection turning moves of these vehicles determine the macroscopic flow splits. In this way, mathematical feasibility (linearization of the macroscopic model) and expressive power (microsimulation of behavior) are combined (Flötteröd and Nagel, 2006a).

FIGURE 1 APPROXIMATELY HERE

Figure 1 shows a simple example. Vehicles move from left to the right. At the diverge they choose from one of three routes, each one having a downstream bottleneck. The figure indicates that the macroscopic density (white=none, green=light, red=jammed) is smoothly synchronized with the vehicles' route choice.

2.2 Behavioral model

FIGURE 2 APPROXIMATELY HERE

The behavioral model requires the availability of an activity plan for every

agent. As an example, consider figure 2a. This plan comprises a three-stage sequence *residence* → *work* → *leisure* → *residence*, which could be typical for an employed person's weekday. Every stage can be associated with a certain activity type and contains at least one location at which this activity can be conducted. In this example, the residence stage is only possible at home, while work can be performed either at the office or at home, assuming that working at home is feasible for this agent. The leisure activity is possible either at home or at a shopping mall.

The agent values the choice of each activity location within every stage according to (a) the direct benefit a choice of this location provides and (b) the expected benefit it can expect from the remainder of its daily plan if it is continued at this location. For example, when comparing the mall and the home location for the leisure stage, a home-working agent has to take into account the cost of traveling to the mall and back home which does not arise if the agent stayed home.

In the figure, the cost of traveling is attached to the links connecting activity locations. The immediate value of choosing a location is expressed in terms of an “immediate cost”, which is taken as the offset-corrected negative of the value. There is also a “remaining cost”, which is the additional cost if afterwards the least cost path is followed. Because of the plan's multi-stage structure, the optimal remaining cost of all activities can be calculated by straightforward dynamic programming, given that the cost of moving through the network is known from the mobility simulation and that the immediate location cost perception is also available from the behavioral model. Overall, this is consistent with a model where travel incurs negative utility while performing an activity incurs positive utility, and travel is only worthwhile if a larger positive utility from the different location overcomes the negative utility of travel.

We generate the multi-stage structure as well as the activity location choice set individually for every agent using output of the MATSim demand modeling and simulation system. MATSim generates a number of alternative activity location sequences, which we combine and reshape such that all alternative sequences fit into a common multi-stage structure (Illenberger, in preparation).

This model allows to effectively represent within-day replanning, as it is clarified in figure 2b. Assume that the agent is about to finish its work stage and leave the office. The choice between going to the mall and going home for leisure can technically be calculated as follows: Add an imaginary destination node to the network and connect all activity locations of the next stage by likewise imaginary links to that destination. Attach the sum of each activity's immediate cost plus its remaining cost to the according link. Then, calculate a time variant best path through the network, with link weights according to the agent's perception of the current traffic situation. The obtained best path does not only yield the subjectively optimal route through the network but also the chosen next activity, which is the last real node in the path.

The possibility to express the combined route and activity location choice by a single best path combination greatly simplifies the behavioral estimation procedure, since it allows to formalize all behavioral issues into a best path problem through a slightly extended traffic network with individualized cost.

3 Traffic state estimation

3.1 Technical description

The estimation method is derived from a Bayesian consideration: We combine a priori knowledge about every traveler's behavior as it is given by its activity plan with anonymous measurements into an a posteriori probability of its behavior given both sources of information. Our method then chooses a route

and a destination for every agent in a way that approximately maximizes this a posteriori probability of the entire population's behavior (Flötteröd and Nagel, 2006b). As stated before, the route and destination choice problem can be subsumed in a single best path calculation, which will be the point of view we adopt in the following discussion.

FIGURE 3 APPROXIMATELY HERE

Some aspects of the simulation system are depicted in figure 3a. It is decomposed into a microscopic representation of traveler behavior and a mixed micro/macro mobility simulation as explained before. In an attempt to realize their individual activity plans, travelers consider their long- and short-term observations of the traffic system state when performing actions within their physical environment. Technically, an agent modifies its path by sending an object representing its perceived cost of traversing the links in the network to a router, which then returns the resulting best path. Note that this cost is individually perceived and can contain perception errors as well as incomplete knowledge.

The behavioral estimation procedure results from reasonable mathematical inference, but can be conveniently illustrated as in figure 3b. Without modifications to the simulation system, the estimation algorithm only modifies the observed travel cost any agent uses to calculate its best path. The resulting behavior is different insofar as it is not optimal with respect to the agent's cost perception any more, but rather with respect to a more general objective function representing the state estimation quality by comparing the simulation output to data from a traffic surveillance system.

3.2 Applicability of the method

3.2.1 Convergent vs. rolling horizon mode

The estimation task can be solved either at once or in a rolling horizon manner.

Since different applications discussed in the next section require different solution methods, a short overview is given here.

In the first case, the estimation procedure iterates over the entire problem time window until convergence. This is called the *convergent mode*. Since a large number of iterations might be necessary and since one needs to know all information (such as measurements) about the entire problem time window a priori, it usually is not amenable to real time operations.

In the second case, only a subinterval of the problem time window is considered at once. This is called the *rolling horizon mode*. By moving this sub-window forwards through time and repeatedly solving the estimation problem only within the sub-window, one usually obtains only a sub-optimal solution to the overall problem. This approach is still attractive if the sub-window's motion is synchronized with real world time such that its solution always provides a most recent estimate of the current system state given the most recent measurements.

3.2.2 Planning vs. telematics application

Transport planning models usually assume that travelers obtained global knowledge of the average system state through many days of exploration and that the resulting behavioral patterns resemble some kind of equilibrium. Typically, such models work at the granularity of an "average day" but do not reproduce within-day fluctuations of the network state as they occur in reality due to the stochastic nature of traffic (Ortuzar and Willumsen, 2004). Estimation in a planning context is possible either in convergent or in rolling horizon mode, where the convergent mode is to be preferred since temporal restrictions on calculation time usually play a minor role in such applications.

Telematics models on the other hand explicitly deal with fluctuations within a specific day. They neither assume global driver knowledge of the system state,

nor do they assume an equilibrium. The behavioral model component in such a system must represent spontaneous driver reactions to fluctuations in the traffic situation (Bottom, 2000). Online applications can only be run in rolling horizon mode since future measurements are not available at the time of the computation. Furthermore, the rolling time window's constant length ensures a constant calculation duration for every iteration.

Both applications can be dealt with by the proposed method. Beyond the chosen operational mode, the only difference is the information about the network state that is repeatedly given to the agents during the iterative estimation procedure. If every agent is provided with global knowledge of the system state as it resulted from the very last iteration, the algorithm converges towards a user equilibrium (UE) in the absence of measurements, and otherwise towards a statistically reasonable compromise between a UE and the observed measurements. If, on the other hand, the information updates given to the agents are constrained to what is actually observable to the agent in that moment, the algorithm does not converge to a strict UE any more, but rather to a solution which realistically regards for a randomly influenced within-day situation. In the latter case, historical information contained in the drivers' activity plans must be used to complement the local observations.

4 Illustrative example

FIGURE 4 APPROXIMATELY HERE

Consider the network shown under various traffic loads in figure 4. Travelers enter the network at the six leftmost horizontal links and leave at the three rightmost horizontal links. Although demand is represented microscopically with each traveler having one origin and one destination, only average occupancies are drawn for readability. The scenario assumes a sensor at the marked link which reports the following velocity measurements: *Free flow speed from 7:45*

to 8:00, low speed from 8:00 to 8:15, again free flow speed from 8:15 to 8:30.

The first column shows the result of a dynamic traffic assignment without incorporation of any measurements. Note that traffic spreads out about symmetrically around the middle horizontal road. Since the sensor information is not available, the method can do no better than doing a time-dependent equilibrium assignment.

The second column shows a result of the estimation procedure described above, *with* the sensor information included. One can observe a traffic jam at the measurement point from 8:00 to 8:15. It can be interpreted as a result of the most plausible overall behavior that resembles available measurements and is consistent with the travelers' original plans.

FIGURE 5 APPROXIMATELY HERE

The underlying calculations are clarified in figure 5. The first column shows relative travel times from the estimation method when the measurement is included. At the marked link, one can see that measurements are qualitatively reproduced. The second column shows the *modifications* of travel times as they result from the estimation procedure that includes the measurement. By replanning based on this modified information, agents do not maximize their subjective utility any more, but the objective a posteriori probability of their actions.

Note that the causalities of traffic flow are properly exploited by the algorithm: The low speeds measured from 8:00 to 8:15 are impossible to reproduce solely by increasing the inflow to the according link, because of its limited flow capacity. The only option is to increase the traffic load on links *downstream* of the measurement location, causing spillback. This effect can be observed in figure 4, right column. The way it is achieved by the algorithm can be seen in figure 5b: Downstream links of the measurement location are made more

attractive by travel time reduction (blue color), thus more driver reroute towards these links and cause the spillback. Since from 8:15 on the measurements indicate free flow, the according cost modifications then are positive in order to keep drivers away from critical links.

Intuitively, this means that the algorithm tests which of the travelers would help best to move the simulation closer to the measurement. This becomes particularly important when multiple sensors are involved, since improving the situation for one sensor may make the situation worse for another sensor.

It is important to note what this algorithm does and what it does not do in response to measurements. What the algorithm does is to modify routes and possibly destinations. What the algorithm does not is to change the traffic dynamics. This means that the algorithm in its current form will *never* estimate an incident (capacity reduction); this is a direct consequence of the modeling assumption that the road network is given and not subject to estimation. The algorithm will instead generate measured traffic congestion from re-routing additional traffic into the congested area. As will be discussed later, it *is* possible and even desirable to combine the algorithm presented here with some kind of incident handling system that handles changes in the physical network.

5 Realistic example

5.1.1 Setting of the test case

We have set up an extensive test case for the proposed method. The geographical zone of investigation is the city of Berlin. Its traffic network is represented by a graph of approximately 2400 links. The MATSim system has been used to generate activity plans for a complete microscopic representation of the Berlin population. The experiments described here use a 10% sample of

this population (approx. 170.000 agents). The network is shown in figure 6.

FIGURE 6 APPROXIMATELY HERE

We gained first experiences with this test case in a real-world application during the soccer world championship 2006. Since we encountered severe problems with all kind of data corruptions (including errors in the network file, unrealistic activity plans, unreliable measurements) during this project, this article considers a setting in which most uncertainties have been removed in order to study the method itself rather than a specific scenario. Accordingly, the results given here are to be understood as a study of algorithmic feasibility. Increasing realism with respect to various sources of disturbances is subject of our ongoing research.

All experiments use synthetically generated measurements as follows: Plans from an imperfect MATSim traffic assignment that did not reach a user equilibrium were loaded onto the network using the same mobility simulation as the estimator itself. For two disjoint 10%-sets of all links, we collected 5-minute averages of the number of vehicles on these links as measurement data. The experiments were run from 6am to 9am, which is the time of the strongest traffic variations in the simulation because of the morning rush hour.

Since the imperfect MATSim result is not a user equilibrium, it can be understood as a behavioral deviation from such, which is exactly the type of situation our method has been designed to handle.

The entire software system is single-threaded and was written in the Java programming language.

5.1.2 Experiments

(a) A priori estimation without measurements

In this setting, the estimator is run without the use of any measurements. As a result, it generates a best assumption of traveler behavior given the MATSim

activity plans by iterating these plans until an approximate user equilibrium is achieved.

FIGURE 7 APPROXIMATELY HERE

The resulting scatterplots are shown in figure 7. In this and all following figures, „link set 0“ and „link set 1“ indicate the measurement subset the according estimation run is compared against. Within each scatterplot, measurement values define the x-coordinate and simulated values the y-coordinate of each point. Accordingly, a perfect measurement reproduction is achieved if all points lie on the main diagonal. The number below each scatterplot indicates its correlation coefficient, which is 1.0 in case of a perfect fit and smaller than 1.0 otherwise.

One observes significant deviation between simulation and estimation. This can be explained by the working of the estimation algorithm in the absence of any measurements: In this case, only the behavioral a priori information is available, which results in a plain user equilibrium assignment as explained above. Since the measurements were generated from a non-equilibrium situation, but are not available to the estimation procedure, the scatterplots represent nothing more but the measurements' deviation from a user equilibrium.

(b) Measurement reproduction

In this setting, parameters were set such that the algorithm attempted to reproduce the measurements by ignoring behavioral a priori assumptions as much as possible: Only measurement-induced cost corrections were visible to replanning agents, while the cost of travel itself was completely ignored. One experiment was run, where the measurements from link set 0 were fed into the simulation. With the resulting estimation, two comparisons were made between estimated and “measured” quantities:

(i) In-sample estimation: Estimated and measured quantities are compared for link set 0, which are the measurements that were fed into the estimation procedure. This comparison tests how much the algorithm follows the measurements at measurement locations. Note that, as is well known, the goal of the algorithm is in general *not* to just follow the measurements, but to follow the measurements as much as is consistent with the model assumptions, one assumption being that also the measurements contain errors.

(ii) Out-of-sample estimation: Estimated and measured quantities are compared for link set 1, in which case none of those measurements were used by the estimation procedure (but the measurements for link set 0 were used instead). This comparison tests how much the algorithm is able to infer, by using its model assumptions, correct quantities at locations where no measurement is available.

FIGURE 8 APPROXIMATELY HERE

The results are shown in figure 8. One observes a very good fit for the in-sample estimation, while the out-of-sample estimation indicates heavy deviation.

The very good measurement reproduction indicates that the method works well. The not *totally* perfect fit is due to various causes, some of which were deliberately accepted while others are still under investigation. Unavoidable but to some degree tunable causes of imprecision are: Incorporation of various mathematical simplifications in the estimation algorithm in order to keep up tractability; use of a random solution mechanism with finite resolution; discretization of all macroscopic quantities in time and space on a quite large scale for reasons of computational performance; use of a linearization based method that might converge only towards a local optimum of the problem.

The heavy deviation between the estimations of one experiment and the

complementary measurement set are not surprising if one remembers that all incentive to reproduce reasonable behavior beyond the available measurements has been removed in this experiment. This is also known as “over-fitting”.

(c) Reasonable combination of information sources

This experiment incorporates the behavioral model with a reasonable weight. As a result, the estimation algorithm abstains from calculating routes that are very unrealistic given an agent's activity plan, even if such behavior yield a better measurement fit. Results are depicted in figure 9.

FIGURE 9 APPROXIMATELY HERE

Again, in-sample estimation as well as out-of-sample estimation are shown.

One notices the following effects:

1. The reproduction quality of measurements involved in the estimation (in-sample estimation) is now slightly worse than in figure 8. The reason for this is the newly incorporated influence of the estimator's behavioral model, which contradicts the unrealistic behavioral nature of the measurements. Since in general measurements are just as error-prone as simulation results, such a compromise is desirable. When judging the estimation quality, it is important to keep in mind that the scatterplots only depict one half of the entire estimation problem; the behavioral fit based on the a priori generated activity plans is not visualized.
2. The reproduction of the complementary measurement set (out-of sample estimation) is now considerably better than in figure 8. This indicates the algorithm's capability to interpolate the traffic state of links that are not directly observed by application of the behavioral model.
3. When compared to figure 7, one observes that the out-of-sample estimation on link set 1 is *not* improved by the inclusion of the measurements from link

set 0. That is, within the current experiment, knowledge of conditions on link set 0 did *not* improve knowledge of conditions of link set 1. The reasons for this are still under investigation since ultimately the method should also yield out-of-sample improvements.

(d) Real time capabilities

Given the general setting as described in experiment (c), the method's real-time capabilities were also investigated. While the previous experiments were run to convergence, here the rolling horizon approach with a time window of 30 minutes was used. The window moved forwards at 1-minute steps, which is approximately the duration of one estimation iteration.

FIGURE 10 APPROXIMATELY HERE

The results are shown in figure 10. Only in-sample estimation results are shown. The rows indicate different estimation step sizes in terms of the agent percentage for which a new path is calculated per iteration. Since a larger percentage implies a faster adjustment but also a loss in precision, the existence of an optimal adjustment rate can be hypothesized.

The results indicate that 2% to 3% is a reasonable value. While the measurement fit clearly increases from 1% to 2%, there is no significant improvement in choosing 3% but at the additional computational cost of adjusting more agents.

This result has an interesting implication. If 2% of the entire population are randomly chosen for adjustment every minute, there still are 30% of the population left after one hour that have not been adjusted at all. While this would definitely become a significant problem if totally unpredictable behavior was to be reproduced, it can be put into perspective by the following two arguments:

1. The incorporation of behavioral a priori knowledge already generates a

„reasonable“ initial assumption of the overall system state, which then is further refined by the estimation procedure. The better the a priori assumptions, the closer are the results to the real state even without use of any measurements.

2. In general there are many route combinations that yield the same network load. From this it can be concluded that for every agent in the population there is a number of other agents with sufficiently similar activity plans to substitute the former agent in an effort to reproduce its observation by a sensor.

Even if much more experiments will be necessary to fully understand all implications of the method, these experiments definitely assert that the algorithm is computationally capable of generating significant estimation improvements in real-time scenarios of realistic size.

6 Conclusions and outlook

We have presented a novel method for behavioral traffic state estimation based on a priori generated activity plans and anonymous traffic measurements. First experiments indicate that the method works with good precision in a real-time setting even for large problems. Still, since the experiments conducted so far only used synthetically generated measurements, many aspects are yet to be explored.

One major simplification was the generation of measurements by the same mobility simulation the estimator itself used. Since model-based assumptions about traffic flow dynamics are currently incorporated as error-free information in the estimation formulation, further investigations with real world data might show that a relaxation of this assumption will be necessary. Since methods for the adjustment of physical traffic flow processes to measurements are available from the systems engineering literature, an integration of both estimation

approaches appears reasonable as stated in the introduction.

A similar statement holds for the occurrence of incidents, which can be considered as structural deviations between modeling assumptions about traffic dynamics and the real situation. The implementation of an additional incident detection module definitely would greatly increase the system's real-world applicability.

An improved a priori demand also implies a better estimation quality. As the experiments have shown, a brute force attempt to only reproduce measurements does not provide a reasonable overall picture of the traffic situation, which makes the incorporation of good behavioral a priori assumptions necessary. This observation suggests a natural operation scheme of the method in a traffic management center: In continuous operations, the estimator could be employed to track within-day fluctuations. If an additional update of the agents' activity plans on a daily basis was realized, the overall system could incrementally improve a transport planning simulation based on these plans as well.

References

- Antoniou, C. (2004) 'On-line Calibration for Dynamic Traffic Assignment', PhD thesis, Massachusetts Institute of Technology
- Ashok, K. (1996) 'Estimation and Prediction of Time-Dependent Origin-Destination Flows', PhD thesis, Massachusetts Institute of Technology
- Ben-Akiva, M., Bierlaire, M., Koutsopoulos, H., and Mishalani, R. 'Dynamit: A simulation-based system for traffic prediction'. DACCORS Short Term Forecasting Workshop, The Netherlands, February 1998
- Bottom, J. A. (2000) 'Consistent Anticipatory Route Guidance'. PhD thesis, Massachusetts Institute of Technology
- Chang, G.L., Mahmassani, H.S., and Herman, R. (1985) 'A macroparticle traffic simulation model to investigate peak-period commuter decision dynamics'. *Transportation Research Record*, 1005:107–120
- DYNAMIT www page, accessed 2005. URL mit.edu/its
- DYNASMART www page, accessed 2005. URL www.dynasmart.com
- Flötteröd, G. and Nagel, K. 'Some practical extensions to the cell transmission model', 8th International IEEE Conference on Intelligent Transportation Systems, Vienna, 2005
- Flötteröd, G. and Nagel, K. 'Simulation and optimization of trajectories in a congested network', Technical Report 06-02, Technical University of Berlin, 2006a. URL www.vsp.tu-berlin.de/publications/
- Flötteröd, G. and Nagel, K. 'Modeling and estimation of combined route and activity location choice', 9th International IEEE Conference on Intelligent Transportation Systems. Toronto, 2006b
- Gingold, R.A. and Monaghan, J.J. (1977) 'Smoothed particle hydrodynamics - theory and application to non-spherical stars', *Royal Astronomical Society, Monthly Notices*, 181: 375–389
- Huisken, G. and van Berkum, E. 'A comparative analysis of short-range travel time prediction methods', 82nd Annual Meeting of the Transportation Research Board, Washington DC, 2003
- Illenberger, J. 'Agent-based modeling of route switching behavior', Technical report, Technical University of Berlin, in preparation
- Kamarianakis, Y. and Prastacos, P. 'Forecasting traffic flow conditions in an urban network: Comparison of multivariate and univariate approaches', 82nd Annual Meeting of the Transportation Research Board, Washington DC, 2003
- Kumar, P.R. and Varaiya, P. (1986) *Stochastic systems. Estimation, identification and adaptive control*, Prentice-Hall.
- Lorkowski, S. and Wagner, P. 'Parameter calibration of traffic models in microscopic online simulations', 84th Annual Meeting of the Transportation Research Board, Washington DC, 2005
- MATSim www page, accessed 2006. URL www.matsim.org
- Ortuzar, J. and Willumsen, L.G. (2004) *Modelling transport*, Wiley, Chichester.
- Schwerdtfeger, T. (1987) 'Makroskopisches Simulationsmodell für Schnellstraßennetze mit Berücksichtigung von Einzelfahrzeugen (DYNEMO)'. PhD thesis, University of Karlsruhe, Germany
- Wang, Y. and Papageorgiou, M. (2005) 'Real-time freeway traffic state estimation based on extended kalman filter: a general approach', *Transportation Research B*, 39:141-167
- Zhou, C. and Nelson, P. 'Predicting traffic congestion using recurrent neural networks', Proceedings of the 9th World Congress on Intelligent Transport Systems, Chicago, 2002

Captions to Illustrations

Figure 1: Screenshot of mobility simulation

Figure 2a: Multi-stage plan structure

Figure 2b: Spontaneous replanning

Figure 3a: Technical overview of simulation

Figure 3b: Technical overview of estimation

Figure 4: State estimation results in a given scenario when a sensor at the location denoted by the arrow reports low speeds from 8:00 to 8:15

Figure 5a: Estimated travel times

Figure 5b: Travel time corrections

Figure 6: Reduced road network of Berlin

Figure 7: Experiment (a): Results of estimation with zero measurement feed.

Figure 8: Experiment (b): Results of estimation in „measurement reproduction“ (over-fitting) mode.

Figure 9: Experiment (c): Results of estimation with reasonable incorporation of behavioral model.

Figure 10: Experiment (d): Real-time capabilities




Figure 1

Screenshot of mobility simulation. Only a cutout is shown in order to increase the figure's resolution.

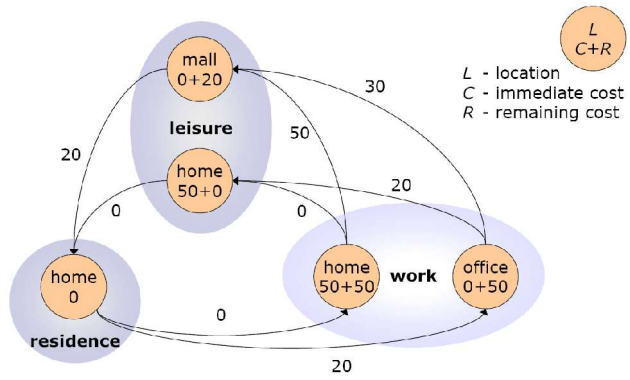


Figure 2a

Multi-stage plan structure.

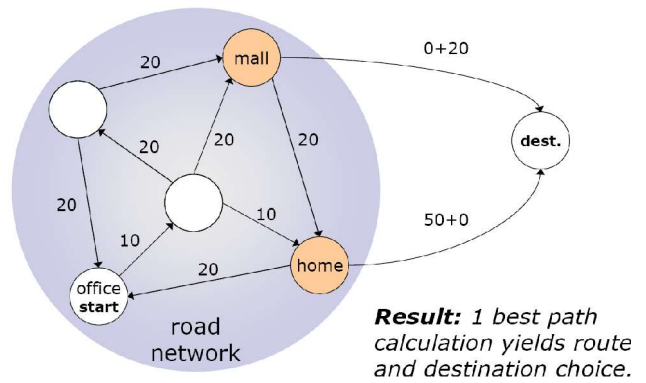


Figure 2b

Spontaneous replanning.

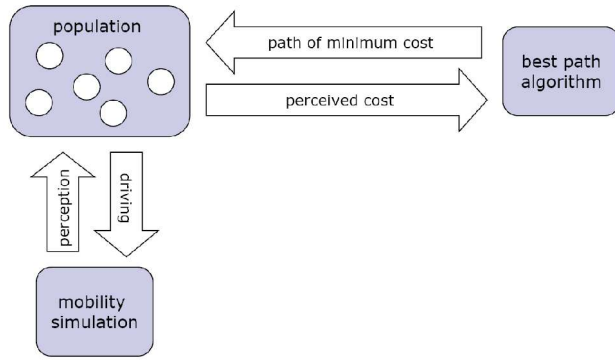


Figure 3a

Technical overview of simulation.

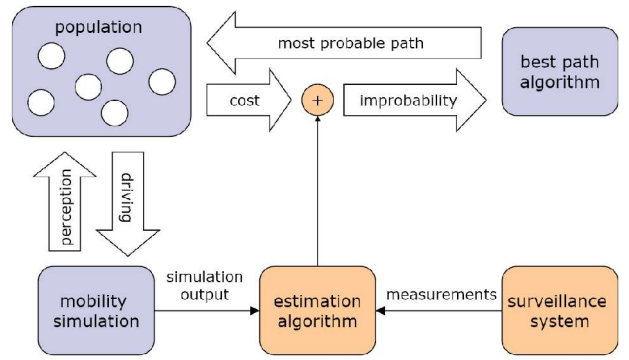


Figure 3b

Technical overview of estimation.

Estimation without measurements

Estimation with measurements

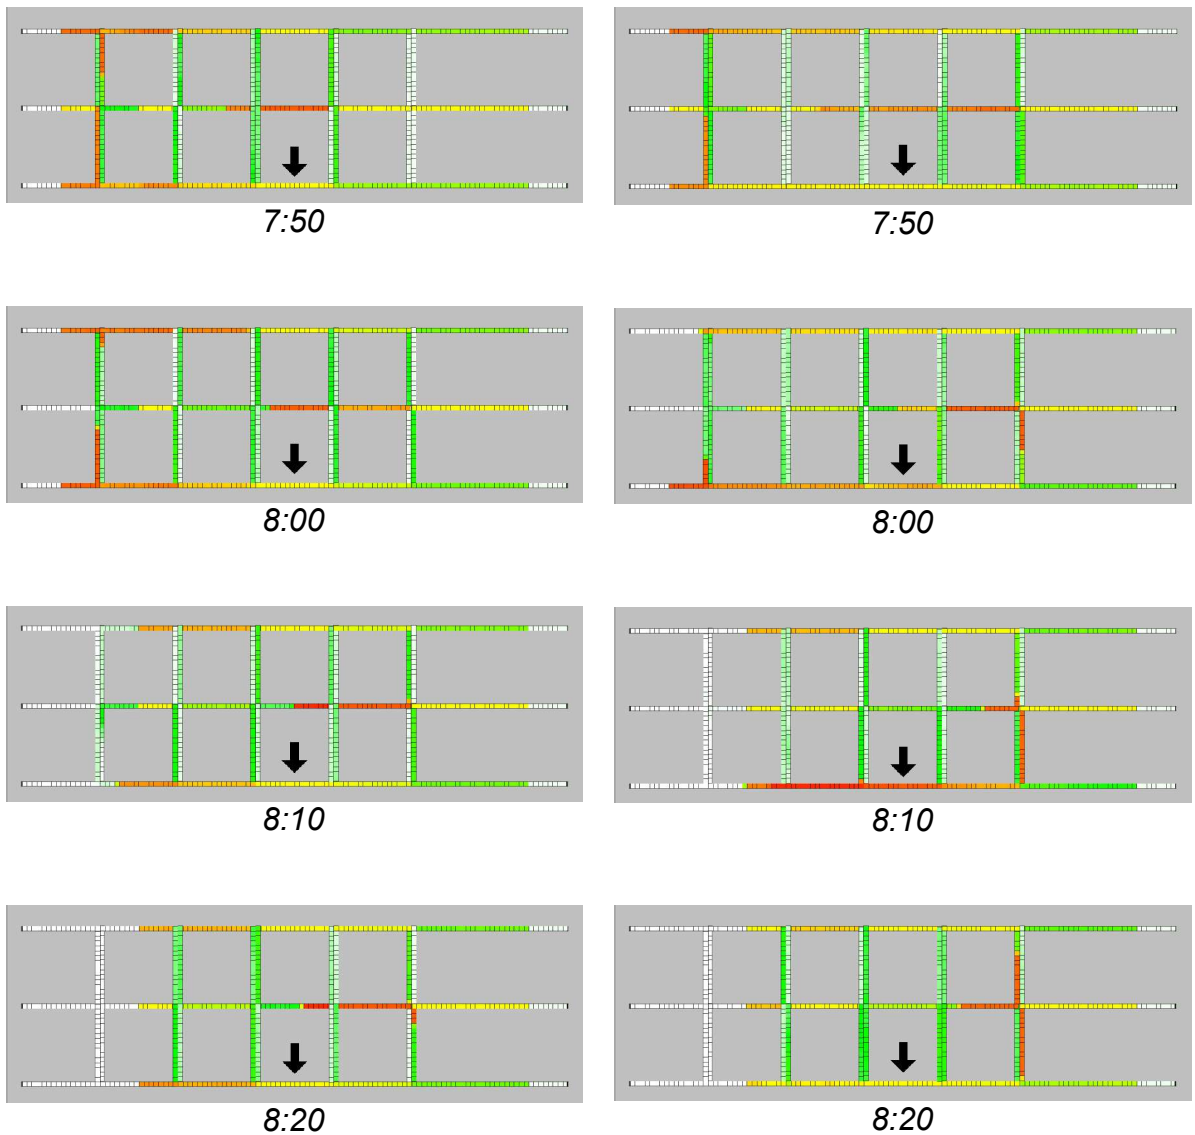


Figure 4

State estimation results in a given scenario when a sensor at the location denoted by the arrow reports low speeds from 8:00 to 8:15. LEFT: State estimation without any sensor input (i.e. the result of the pure dynamic traffic assignment). RIGHT: State estimation when the sensor information is used. – Colors indicate relative road occupancy: white=empty, green=light, yellow=high, red=jammed.

Estimated travel times







Figure 5a

Colors indicate actual travel times:
white=minimal; green, yellow=
increased; red=considerably increased.

Travel time corrections







Figure 5b

Colors indicate modifications of travel
times: blue=negative; white=zero;
green, yellow, red=increasingly positive.



Figure 6

Reduced road network of Berlin.

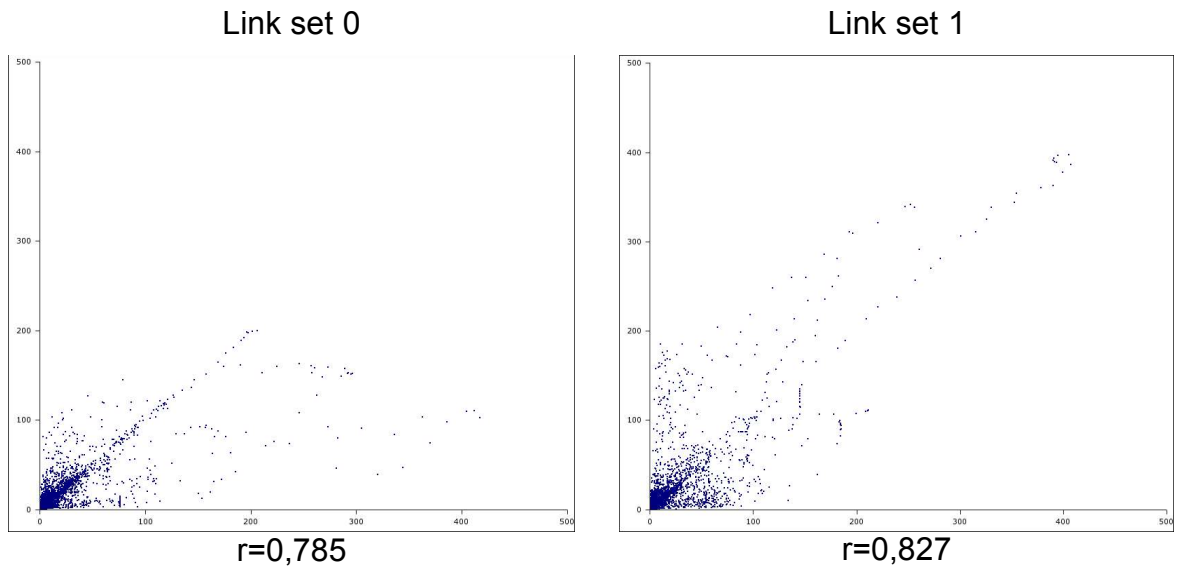



Figure 7


Experiment (a): Results of estimation with zero measurement feed.

In-sample estimation (link set 0)



$r=0,988$

Out-of-sample estimation (link set 1)

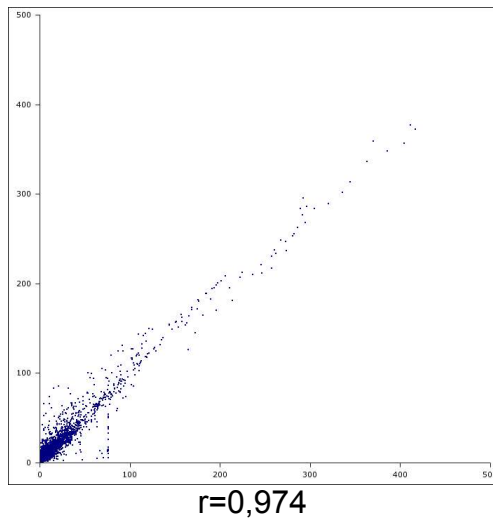


$r=0,445$

Figure 8

Experiment (b): Results of estimation in „measurement reproduction“ (over-fitting) mode.

In-sample estimation (link set 0)



Out-of-sample estimation (link set 1)





Figure 9


Experiment (c): Results of estimation with reasonable incorporation of behavioral model.

In-sample estimation (link set 0)

1%



2%



3%




Figure 10

Experiment (d): Real-time capabilities.