The Knowledge Engineering Review, Vol. 29:3, 375–403.
doi:10.1017/S0269888913000118

A review on agent-based technology for traffic and transportation

ANA L. C. BAZZAN[1] and FRANZISKA KLÜGL[2]

[1]*Instituto de Informática/PPGC, UFRGS, Caixa Postal 15064 91.501-970 Porto Alegre, RS, Brazil;*
e-mail: bazzan@inf.ufrgs.br;
[2]*Örebro University, Fakultetsgatan 1, 70182 Örebro, Sweden;*
e-mail: franziska.klugl@oru.se

Abstract

In the last few years, the number of papers devoted to applications of agent-based technologies to traffic and transportation engineering has grown enormously. Thus, it seems to be the appropriate time to shed light over the achievements of the last decade, on the questions that have been successfully addressed, as well as on remaining challenging issues. In the present paper, we review the literature related to the areas of agent-based traffic modelling and simulation, and agent-based traffic control and management. Later we discuss and summarize the main achievements and the challenges.

1 Agents and multiagent systems: promises to traffic and transportation

The increasing demand for mobility in the 21st century poses a challenge to researchers from several fields. More efficient tools and techniques are necessary in order to deal with traffic and transportation problems, including control, optimized use of the existing network, efficient assignment of the demand, and so on. More than ever, interdisciplinary approaches are necessary.

In particular, a successful experience has been the cross-fertilization between traffic, transportation, and artificial intelligence that dates at least back to the 1980s and 1990s. During the last decade, there has been a tremendous progress in transportation and traffic engineering due to the use of agent technology. Given the increasing complexity of transportation and traffic systems, which arises from the modern way of life and new means and organization of transportation, not only new techniques must be deployed, but also the individual choices must be better understood if the whole system is to become more efficient. Thus, it is not surprising that there is a growing debate about how to model and improve traffic and transportation systems at both the individual (micro) and the society (macro) level. This may raise technical problems, as transportation systems can contain millions of autonomous, interacting intelligent entities that need to be simulated and, perhaps, controlled. Therefore, traffic and transportation scenarios are extraordinarily appealing and challenging for (multi-)agent technology. Additionally, traffic scenarios have become very prominent as test beds for coordination and adaptation mechanisms in multiagent systems.

In summary, it is well established that agent-based approaches suit traffic and transportation management very well given the geographical, functional, and temporal distribution of data and control, as well as the frequent and flexible interaction between the participants and their environment. Therefore, agent-based approaches can contribute to the whole effort around the design and control of intelligent transportation systems (ITS) and ultimately to make our cities indeed smart. In particular, one may identify a number of main motivations for using agents and multiagent system technologies in traffic and transportation:

- Natural and intuitive problem solving by active entities with a (potential) local perspective, instead of complex, central solutions for which the inclusion of all necessary details and constraints is not feasible. Adaptive and robust services can be provided due to their self-organization capability.
- Autonomous agents provide an appropriate basis for modelling heterogeneous systems. Every entity may possess its individual architecture, state representation, and behaviour. Thus, an arbitrary level of detail can be included into a simulation model or an arbitrarily sophisticated problem-solving framework can be applied on the agent level. The integration of legacy software is facilitated by using an 'agent wrapper'.
- Agents and their interaction can be described using high-level abstractions. Thus, they provide an intuitive level of interaction between human users or modellers and the agent-based system. Here, an important related issue is the visualization, which facilitates the analysis and control of microscopic properties of a given system.
- Agents or multiagent systems technologies allow coping with variable structure of the system in an elegant and efficient way. If the active entities are modelled as agents—in control and management applications as well as in developing simulation models—they may control when and with whom they are interacting. These (dynamic) relations may be controlled, disconnected, or established from the local point of view. Agents may be able to adapt their behaviour to a changing organization. This flexibility is highly relevant in traffic and transportation domains.
- The agent metaphor used for modelling a traffic participant or decision-maker enables us to capture complex constraints connecting all problem-solving phases. The agents (and their reasoning capabilities) may be persistent in their context and environment. This leads to the possibility of tackling entities over their complete 'life span' in a consistent and coherent way.

To some extent, agent and multiagent technologies are well understood. Solutions have been suggested for many problems and scenarios related to transportation and traffic engineering such as: reproducing human behaviour, pedestrian and vehicular flow simulation, distributed optimization and control, management of diverse kinds and combinations of systems (public transportation, traffic, air traffic, etc.), and various levels of decision making.

This scientific and practical progress is reflected in the large number of papers dedicated to these and related subjects. In particular, the international workshop series on 'Agents in Traffic and Transportation' (ATT) provides a discussion forum for researchers and practitioners from the fields of artificial intelligence—in particular autonomous agents and multiagent systems—and transportation engineering. Therefore, in the remainder of this article we take a deeper look into the achievements of the last decade, into the issues that were already addressed as well as into the challenges and opportunities for further research.

This text is organized in two large areas:

1. Agent-based traffic and transportation simulation (Section 3). This aims at reproducing human decision making and behaviour, capturing the level of detail that is necessary for a particular objective. Modelling and simulations here aim at supporting the analysis of the effects of decision making in traffic-related contexts at various levels of abstraction. Examples are the analysis of abstract scenarios contrasting user and system goals, and co-adaptation for predicting the effects of given measures and strategies in a particular network. This part is organized around different levels of human decision making in traffic scenarios. Thus, we start with the demand side of a traffic and transportation system in terms of agent-based travel demand model, then continue to discrete decisions mainly for route choice and finally end with the lowest level, namely, simulation of actual driving.
2. Agent-based management and control (Section 4): we introduce and discuss concepts of distributed, yet coordinated decision making by intelligent entities in order to develop more robust and efficient solutions for controlling traffic flow and/or for managing transportation systems.

This part starts with optimization-based approaches of single and networked controllers, then continues with learning controllers (again, single and network cases). This part is followed by approaches that take the demand into account.

Before these two sections, in Section 2, we give a brief introduction to the most relevant concepts in agents and multiagent systems. In Section 5, we summarize the achievements so far and the remaining challenges related to the use of agents and multiagent systems in traffic and transportation. Concluding remarks appear in Section 6.

At this point, we remark that there is a number of other related topics that could be part of a broader overview: simulation of pedestrian movement, public transportation, or new trends in agent-based freight simulation as well as logistics, air traffic control, railway management, and innovative topics such as cooperative driving, platoon formation, or agent-based approaches to simulation of networks where vehicle-to-vehicle (V2V) communication exist. Readers interested on a review about agent-based management of logistics are referred to the overview provided by Davidsson *et al.* (2005). A recent short and relatively technology-centered review by Chen and Cheng (2010) gives pointers to a variety of these topics. Due to the restriction of space, and with the intention of being more focused and analytic, in the present text we restrict ourselves to traffic simulation and control as previously outlined.

2 Basic concepts of agent and multiagent system technology

One reason for the popularity of agents and multiagent systems relates to the advances in computational systems, which are increasingly more distributed, open, large, and heterogeneous. The positive side about this is visible every time we use the Internet. However, this comes at a cost: increasing complexity of both hardware and applications. Canonical examples are the current paradigm of cloud computing versus centralized, mainframe-based computation of some decades ago.

Managing interactions among autonomous entities with ever increasing interdependencies has been one of the biggest motivations for distributed artificial intelligence and for multiagent systems. These aim at developing and analysing models derived from social interactions in human societies, and applying them to computational systems in order to resolve conflicts in organizational structures of various types.

In particular, in the present paper we deal with agent-based techniques for modelling and simulation as well as for coordination and adaptation. As a comprehensive introduction to the various related subjects is outside the scope of this paper, we simply sketch the basic ideas.

Agent-based simulation uses the metaphor of autonomous agents and multiagent systems as the basic model conceptualization. This means that a model consists of interacting agents situated in a simulated environment. Agents may correspond to cities, blocks, platoons, households, individual travellers (drivers), vehicles, sensors, traffic signals, etc. Also elements of the environment may be conceived as agents. A general introduction into the underlying concepts of agent-based simulation from the social science point of view, together with a few selected examples can be found in Epstein (2007) and Gilbert (2007) gives a practical introduction to the basic ideas, model development and its implementation; a more recent introduction to more general aspects of agent-based modelling and simulation can be found in Klügl and Bazzan (2012).

In the context of simulation modelling, two concepts are particularly relevant: BDI (*Belief–Desire–Intention*) and layered architectures. In the particular case of traffic modelling, both are useful for combining decision making on different (e.g. temporal) levels such as high-level strategic planning (e.g.route choice) and low-level tactical action selection (e.g. actual driving).

3 Agent-based traffic simulation

Traffic simulation represents a prominent application for modelling and simulation. It supports complex urban and transport planning, as well as management tasks on different levels of granularity in space and time. It is one of the prerequisites for testbeds for modern technology such as intelligent traffic control or V2V communication systems. Also, simulation can be used for generating background traffic in driving simulators that are used for training purposes. Some objectives of traffic simulation studies are: to improve the understanding about how decisions by individuals lead to overall traffic phenomena (e.g. dynamic congestion, emissions, etc.); analysis of the effect of changes in infrastructure or in control strategies in specific scenarios; testing the feasibility of particular communication forms; or producing short-term extrapolations of the current traffic state.

During the last years, agent-based approaches to traffic simulation have shown that they are able to capture necessary details at entity level as well as to reproduce relevant realistic phenomena. Agents may represent drivers, vehicles, or other traffic participants. They are explicitly present as active, heterogeneous entities in an environment representing the road network where they may exhibit arbitrary complex information processing and decision making. Their behaviour, especially those resulting in simulated movement, can be visualized, monitored, and validated at individual level, leading to new possibilities for analysing, debugging, and illustrating traffic phenomena. Restricting factors are computational power, computer memory, and—the most critical aspect—available data.

After a short introduction to the motivation for agent-based approaches to traffic simulation, we review the state-of-the-art in this area. We do this from larger time frames to smaller ones: we start with approaches for generating travel demand, continue with models that aim at capturing route and other traffic-related choices, and finally to particular agent-based traffic flow simulators. Also, simulations that integrate these three levels of agents' decision making in traffic and transportation are described.

3.1 Why agent-based traffic simulation

Traffic and transportation are attractive application domains for agent-based simulation for several reasons: for example, mobility is a highly relevant topic for societies. Hence, any improvement achieved using a new technology has effect on the daily life of many people as well as on the economy. Thus, traffic-related topics are on the agenda of many funding agencies. Often, traffic management and control approaches require traffic simulation for illustrating functionality or performance tests. On the other hand, traffic scenarios are ideal testbeds for new (general) algorithms as they are easily explainable, given that everybody has to deal with several aspects of mobility in his or her daily life. Also, dealing with traffic simulation is interesting from a methodological point of view due to many challenges ranging from questions on necessary level of detail and required knowledge, the presence of emergent phenomena, spatial distribution, or the integration of different goals present in traffic systems. Last, but not least, the scale of the system that needs to be treated is a challenge on its own.

Using agents as the basic metaphor for reproducing intelligent decision-making entities offers many advantages compared with conventional approaches in traffic simulation. First, it allows an elegant treatment of heterogeneous and variable structures in the agent population (e.g. individual drivers entering and leaving the system at given time periods), as well as in the transportation network (links may be deleted or created during a simulation and agents adapt to these changes). Second, it enables the modelling of complex information processing and decision making considering multiple factors and dynamic information (including anticipation, group behaviour, or learning and adaptive behaviour). An agent-based approach enables an intuitive integration of behavioural constraints throughout different levels and phases of the decision-making process. For example, in simulations related to multi-modal routing, a decision made by a given agent about taking a bus has consequences on further decisions of this agent. One consequence is that the option of using a car is no longer available for this particular agent in future decisions. Third, since every active entity is explicitly represented, they not only support distinct and elaborated visualization, but also permit modelling heterogeneity on different levels (e.g. based on different contexts, different parameter sets or even different architectures).

3.2 Alternative microscopic approaches to traffic simulation

The just mentioned features distinguish agent-based approaches to traffic simulation from more conventional microscopic ones. In the four-step model (Ortúzar & Willumsen, 2001), the initial phases basically consist of data collection and manipulation for determining the number of trips between locations, followed by the assignment phase, which connects a trip to a route resulting in the simulated load on the network. Further, depending on the objective of the study, a mobility simulation is used for determining the actual traffic flow data. Agent-based approaches can be used not only for each of these phases, but also to emulate the decision making across all involved aspects. For example, instead of identifying flows between areas where a significant share of the population resides and industrial areas for determining demand for transportation, an agent-based approach uses models for explicit individual decisions about why and where an agent wants to be. This means individual context-dependent decisions on locations, times, and activities. The most prominent methodology for modelling agent decision making is based on econometrics de Palma *et al.* (2008). Here, weights and parameters of an individual function for evaluating different alternatives in decision making are optimized for minimising the gap between measured (stated) and simulated distribution. The main outcome are the weights that describe preferences of the traffic participants. Some issues here refer to the required quality of input data, the limited complexity of the functions that can be estimated, as well as the transferability of results to similar scenarios.

Game-theoretic analysis can be seen as an alternative microscopic tool providing means for the analysis of interaction when the overall outcome depends on all individual decisions. Interactions are formulated as an abstract game, whose analysis is based on payoff matrices or functions describing the possible outcomes for the different participants. This allows the identification of equilibrium states that may predict the selection of rational actors. Examples of such methodology appear in Section 3.4: Using game-theoretic metaphors such as minority or congestion game, agents choose one alternative. Those that have selected the less-populated option win or receive the highest reward. When played repeatedly, those games provide a good abstract model for many traffic scenarios in which agents want to avoid selecting the alternative that is selected by the majority. Game-theoretic approaches in traffic scenarios cannot be separated from agent-based approaches. In particular, abstract scenarios such as the minority game are mostly tackled within the realm of agent-based simulation. Yet, in general, agent-based simulation is more powerful as also 'irrational' behaviour, decisions involving a lot of information, dependencies to other agents, etc. can be integrated into the model.

In traffic flow simulation, one can find a number of microscopic techniques. Highly elaborated car-following models involving vehicle dynamics and driver decision making such as the model by Gipps (1981) are used in commercial tools (see Barceló, 2010) for a collection of contributions on various tools based on driving strategy models). Such models basically resemble elaborated automated cruise control algorithms for simulated vehicles with individual values for desired speed and acceleration or acceptable gaps to the preceding car. These models can be easily conceptualized as simple agent-based models. Yet, agent-based models can not only integrate higher level of decision making, but also more complex forms of information processing, which may involve more than headway gap and speed of the car ahead.

Another microscopic traffic simulations that aim at reproducing traffic flow are based on cellular automata. Space is segmented into cells whose state is updated based on the state of their neighbours following a set of update rules. A well-known model is described in Nagel and Schreckenberg (1992), where the state of a cell represents the presence and speed of a vehicle on that cell. The positions of vehicles are updated based on four rules that express speed adaptation with respect to the gap between the agent and its preceding vehicle, and on a probabilistic deceleration rate.

It is interesting that cellular automata model are often described from an agent-based point of view, characterizing the behaviour of an active vehicle–driver unit instead of a spatial unit. Cellular automata can be implemented very efficiently, thus making them suitable for large-scale traffic flow simulations. However, the necessary complexity of the decision making captured in the state of a cell and the formulation of strictly local update rules can be quite challenging when involving more than the simple rules of the Nagel–Schreckenberg model.

Queueing simulation is another efficient technique for traffic flow simulation that abstracts the actual decision making during driving on a link by capturing the time that a vehicle–driver unit stays on a particular link based on its position in a queue. For an example, see Charypar *et al.* (2007). Queueing simulation can be seen as an intermediate solution between microscopic and macroscopic simulation paradigms. Vehicle–driver units are passively moved through the system, that is, there is no active decision making while driving. Nevertheless, it is very efficient and fits well an agent-based approach for other levels of decision making. The agent-based traffic simulator MATSim (`www.matsim.org`) uses a queueing approach for simulating the actual movement combined with an agent-based demand generation and departure time, mode, and route choice, as described in the next section.

3.3 Agent-based travel demand simulation

Traditionally, determining the travel demand is the initial phase in all traffic simulation. It deals with the need for transportation infrastructure as the basic simulation input. The output of travel demand calculations is the number of trips from an origin to a destination, eventually also considering departure times. Determining travel demand is traditionally a data-driven activity based on demographic data, statistics about workplaces and households, car ownership, etc. Whereas trip-based approaches were prevailing in the past, activity-based approaches are becoming more and more important. In these, the daily schedule of typical human that belongs to a particular behaviour class is reproduced. It consists of activities that happen at particular locations. In order to simulate changing in locations, trips are added to the overall plan. After this, origin, destination, and also departure time can be determined based on the respective trip connecting two activities (Bhat & Koppelman, 2003; Timmermans, 2005).

Agents and activity-based approaches perfectly fit together. Simulating activities planned and executed by agents enables the analysis of traffic systems on a more detailed level and thus a wider range of transportation policies:

- Increasing complexity related to socio-economic properties and other sources of heterogeneity can be handled as decision-makers and their properties are explicitly represented.
- Individual adaptation and learning of daily plans can be used for acquiring realistic and optimized plans.
- Integration of interactions between agents can be included: for example, using social network models, but also between members of the same household.
- Integration of other levels of decision making such as mode and route choice as well as actual driving and potential delays gives immediate feedback on the feasibility of the (remaining) daily plan.

Therefore, it is not surprising that a number of agent-based concepts and implementations for travel demand models exist.

Rindsfüser *et al.* (2004) propose a model of an intelligent agent for adapting daily activity schedule with respect to external events. The starting point is the definition of an habitual daily program as a coarse pattern for a traveller's activity schedule. This is extended and adapted over the simulated day in reaction to the traffic situation, meetings with friends, etc. This 're-scheduling' model illustrates the necessity of flexible human decision making for producing realistic daily plans. Although a first prototype could produce realistic plans, the agent model was not applied in a large, realistic setting.

Auld and Mohammadian (2012) separate activity generation from activity planning (configuration) and scheduling (the actual execution during simulation) into a dynamic decision-making framework related to activity and travel planning. The model is grounded on empirical studies. Similarly, Sun *et al.* (2012) provide additional grounding dealing with flexible information handling and uncertainties in decision making.

In contrast to this, the activity-based travel demand module of the aforementioned MATSim—the so-called MATSim-T module (Balmer *et al.*, 2009)—does not try to resemble the planning process. Rather, it is based on a set of fully elaborated daily plans that are iteratively adapted. This makes an activity-based approach feasible for large-scale cases, such as a simulation of the complete Zürich area. In MATSim-T, each agent possesses a set of fully elaborated daily schedules, including details of the route choice to connect two (or more) locations. Schedules of all agents are simulated using a queueing simulation for actually producing realistic simulated travel times on the road network. Travel times and costs are then evaluated using a fitness measure and optimized by means of genetic algorithms for adapting the plans. In a next iteration, the agent uses the evaluations to select one of the plans for execution. Modules are available, for example, for the selection of destinations (Horni *et al.*, 2009, 2011). The optimal schedule is then taken as the most realistic one.

Grether *et al.* (2008) have used MATSim-T to illustrate the usefulness of an agent-based simulation to evaluate time-dependent tolls in road pricing. They were able to show the effects of an afternoon toll in some roads in central Zürich. Such a detailed analysis is only possible if the simulation is able to capture decisions at different times of a day, as well as the different activities and travel needs of agents that can adapt their complete daily plan to avoid (or not) certain areas at certain times.

The idea of agent- and activity-based demand generation coupled to traffic simulation was also transferred to the simulation of commercial vehicles delivering goods in an urban environment (Joubert *et al.*, 2010).

The research team around T. Arentze and H. Timmermans have developed activity-based models for travel demand generation grounded on existing theories in psychology and economics. In Arentze and Timmermans (2008), the focus is on the allocation of tasks and budget within households for elaborating the activity schedule at household level. In Arentze and Timmermans (2005), they propose the use of Bayesian networks and their adaptation for developing a model of mental maps as an individual representation of the user's environment. These mental maps typically contain incomplete and incorrect information. Using a form of cognitive learning, the authors aimed at overcoming drawbacks in previous works on activity-based travel demand, in which choice sets were often specified in an *ad hoc* way. Their approach was then extended in Han *et al.* (2009), here dealing with dynamic location choice sets based on a combination of different forms of learning and psychological theories such as Bayesian learning or social comparison theories.

Activities are not planned or scheduled in isolation: social interactions are relevant for joint activities, coordinated activities, social influence on preferences or information exchange. Ettema *et al.* (2011) illustrate the relevance of social interactions in activity-based simulation; Hackney and Marchal (2011) describe a tool that allows the analysis of the effect of social connectivity and location on activity and travel choices. Ronald *et al.* (2011) analyse the influence of different interaction protocols on number, frequency, and type of activities when more than one simulated individual is deciding about an activity. Also the influence of social networks on destination choice and thus on traffic demand has been tested, as in Han *et al.* (2011) and Horni *et al.* (2011).

An earlier approach to activity-based travel demand modelling was suggested by Rindt *et al.* (2002): the Contract Net Protocol was used for negotiation about participation in an activity. Unfortunately, the paper does not provide enough details to assess how this highly dynamic activity generation process could be integrated into a larger traffic simulation process; nevertheless, it gives an interesting bottom-up alternative to other highly refined approaches.

A non-activity-based approach for travel demand modelling is described in Zhang and Levinson (2004). Their model generates travel demand from interaction between travellers, node, and link agents. Travellers aim at finding and reaching an activity that is associated to a node agent. Link agents manage a set of reachable activities together with shortest paths and costs for connecting them. Thus, this model combines trip generation with route assignment, rather than dividing them as in conventional approaches. As the previous model, this approach illustrates what can be done directly based on the agent concept. It remains to be tested whether it produces a valid output when applied to real-world settings.

One of the decisions that is usually tackled in the context of travel demand modelling is departure time choice. The relation between this choice and congestion is analysed in McBreen *et al.* (2006). Special attention is paid to the influence of homogeneity or heterogeneity of agents decision making on the overall dynamics. Their results indicate that only models that are capable of considering heterogeneous agents can produce robust outcomes.

3.4 Agent-based simulation of traffic-related choice

A second group of agent-based models deals with traffic-related choice. These choices may be made on different levels of granularity: from destination and departure time choice, mode, and route choice, to choice related to lane change. In such models, the agents (repeatedly) select an option out of a set of alternatives. The result of these decision processes are particular outcomes or utilities for the decision maker. For sake of clarity and because it is an important choice dimension, we focus on route choice. We remark, however, that route and mode choice can be integrated and performed simultaneously.

The research area of discrete choice analysis deals with such choice processes, by adapting micro-level decision functions or weights from a macro-level point of view, aligning empirical survey data and stochastic choice results. In contrast to this, an agent-based approach is usually bottom-up, starting from explicit individual, context-depending behaviours. Agent-based approaches are particularly relevant when networks are dynamic, when dynamic information is available, or when there is not enough specific data available for using an econometric approach. In this case, an agent-based simulation may integrate the agents' immediate reaction to the changed environment (including other drivers). Also, complex information from different sources can be included into route choice decision making as the agent-based modelling allows integrating sophisticated decision making with individual preferences or socio-economic background.

3.4.1 Game-theoretic route choice models and information

Two main directions can be found in the literature. In one, the overall route choice scenario is abstracted and techniques from game theory are used. In the other, scenarios with more than just a few route alternatives, often using a realistic road network, are explicitly represented and more complex decision-making models are used.

Regarding the former, abstract scenarios are mostly inspired by congestion or minority games. The basic idea is that agents decide simultaneously between two routes; those that select the less crowded one receive higher reward. Repetition of decision making is coupled to some adaptation or learning strategy so that agents' choice is adapted to the feedback received in previous iterations. This way, an user equilibrium may be reached (Wardrop, 1952). In terms of game theory, this means that no agent can improve its reward or reduce its costs by switching routes without worsening at least another agent (Pareto optimality). Examples of such abstract two-route scenarios can be found in Chmura and Pitz (2007) or in Klügl and Bazzan (2004). In both, a reinforcement learning scheme is used for reproducing the decision making of human subjects in a corresponding experimental study. In the latter, an additional phase is introduced for allowing all or some of the agents to re-consider their choice based on a forecast produced from the original set of choices made by the agents. Thus, the agents not only learn about selecting a route in the first phase, but also how to evaluate the information received in the second phase. One important result of this study was that a certain share of agents that ignore traffic forecast turned out to be necessary for an efficient agent adaptation. Learning and information sharing for coordination of route choice in an abstract two-route scenario is also the topic of Machado and Bazzan (2011). Here, random Boolean networks are used to integrate the information from other drivers into the decision making of each agent. This way, it was possible to coordinate agents' choices to the optimal distribution among both routes. Finally, there are approaches that are not based on learning. For example, Bazzan and Junges (2006) have shown how congestion tolls can serve as information for guiding the bottom-up coordination of agents in an abstract two-route scenario.

In these game-theoretic scenarios, the reward of agents when selecting a route is calculated based on the number of all agents that chose the alternative. This is a very abstract view and does not resemble the actual dynamics of traffic situations. Objectives such as finding the appropriate form of information that leads to equilibrium states cannot be accomplished with such abstract scenarios. Nevertheless, for basic analysis of decision-making dynamics, abstract scenarios are satisfactory as they allow focusing on the really relevant aspects.

Apart from abstract scenarios based on game theory, other techniques are used for analysis of information effects in two-route scenarios. In order to analyse the effect of specific information types on route choice, Wahle *et al.* (2002) use a microscopic simulation of the traffic flow based on the Nagel–Schreckenberg rules. Information about the traffic state is communicated to the drivers in order to support their route choice. This information is generated by floating cars that transmit their individual travel times after having finished their trips. In this scenario, the authors analysed the effect of different quality of information. In some settings (e.g. giving delayed information to drivers), harmful oscillations occur regarding the numbers of drivers on each route. In subsequent studies results were presented, which dealt with using other forms of information calculated from a macroscopic point of view, such as overall density, mean velocity, or information regarding trends. These ideas were further elaborated in Klügl *et al.* (2003) where agents could learn which information was the most useful. In Bazzan and Klügl (2005), it was shown that providing information can be also useful in the context of the Braess paradox.

3.4.2 Models involving network learning, communication, and re-routing

Although the study of Yamashita *et al.* (2004) did not aim at reproducing a particular real-world system, they have used simulation for testing the effect of traffic information (see also Section 4.9 where we discuss the issue of providing information for control purposes). Their concept of 'route information sharing' assumes that drivers communicate details about their planned route to a server that broadcasts aggregated route information. They tested the system in various network topologies such as a grid or a radial ring structure. Their results showed that an increasing share of informed drivers yields decreasing travel times not only for drivers equipped with the information sharing device but also for the others. Moreover, travel times of drivers equipped with this device were lower than that of drivers without it. Later, this idea of sharing route information was further developed into a collaborative navigation system (Yamashita & Kurumatani, 2009). Also, Zhu *et al.* (2008) were able to show that an agent-based approach combining route information exchange with random walk in a road network (without cycles and U-turns) can produce valid results (in comparison with existing discrete route choice models), while being still computationally tractable.

In Klügl and Rindsfüser (2011), simulated drivers repeatedly learn about realistic travel times on the used network segments for producing a balanced distribution of load in a real-world network. After, a link is blocked and drivers close to this link are informed about this blockage. Agents are then allowed to re-route. It was shown that the load distribution in this situation was quite different from the situation in which the blocked link was known from the beginning. The value of re-planning for drivers with different information settings was investigated in the work of do Amarante and Bazzan (2012).

The effect of real-time information on drivers' decision making was also studied by Dia and colleagues. In Panwei and Dia (2006), a fuzzy neural architecture was used, where socio-economic parameters were represented as fuzzy variables for simulating decision making about whether or not to keep their initial route decision when new information was made available. Simulation results were empirically tested and validated. In Dia (2002), survey data gathered from commuters in a congested area of Brisbane (Australia) was used to analyse relevant factors using an econometric analysis. The results were integrated into an agent-based simulation with a BDI agent architecture. Thus, the author could demonstrate the feasibility and usefulness of such complex agent architectures in a practical use. In the next subsection, more examples of the use of similar agent architectures for decision making in traffic scenarios are given.

3.4.3 Complex agent architectures for traffic-related choice

The works discussed in the previous sections are mostly based on simple agent architectures, where the agents' behaviour is determined using, for example, a restricted number of rules in a reactive way. However, more complex agent architectures have been proposed for reproducing more elaborate decision making, as discussed next.

In Rossetti *et al.* (2002), an extension of the DRACULA traffic simulator is described. It combines demand generation, departure time choice, route choice, and actual micro-simulation of the movement with an agent-based formal model for the cognitive decision making of the drivers. BDI-based agent specification and Agentspeak(L) were used for modelling decision making without information, with information given only before the start of the travel, and with information given both before and during the trip. Agents update their belief base and select plans (routes) and departure time.

Layered agent architectures for simulated travellers was proposed in Bazzan *et al.* (1999). It combines BDI-based high-level reasoning concerning route choice (strategical layer), with a reactive layer for actual driving (tactical layer). A similar organization into two layers of agent behaviour—the physical mobility and the mental decision-making layer—can also be found in the work by Balmer *et al.* (2004).

3.5 Agent-based traffic flow simulation

Depending on the objective of the investigation under course, simulation of route and mode choice is not sufficient. In some cases it may even be irrelevant. Most of the previously discussed models deal with time granularity of, say, 1 day. These are useful for, for example, calculating travel times. On the other hand, the traffic flow simulations discussed next operate with a time granularity of 1 minute or less. There are several possible applications for such models: First, a valid reproduction of traffic flow can directly address traffic phenomena like the emergence of traffic jams based on driving behaviour, showing the effects of or speed limits, among others. Other examples are the prediction of traffic situations or interpolation of the current traffic state between distant located sensors. As every vehicle is treated as a distinct entity, individual-based emission of pollutants can be calculated if data related to a heterogeneous population of vehicles are available. Finally, agent-based simulation of traffic flow is very attractive for visualization of details, also generating background traffic for (immersive) traffic simulators in which a human subject virtually participates in traffic. There are even games as, for example, 'Cities In Motion'[1] using this level of traffic simulation.

This level of traffic simulation already is and will continue to be highly relevant for testing future communication infrastructures such as V2X (a generalization of V2V) communication and *ad hoc* networks, future infrastructures for safe and informed driving, etc. Many of the approaches discussed in the next section reproduce, as realistically as possible, situations such as how simulated drivers approach a traffic signal. Clearly, although sometimes relevant, routing is not the focus of this research. In the following, we tackle only those simulations studies with a focus on driver behaviour. Further examples of the use of the mentioned communication infrastructure in conjunctions with simulation based on new technologies appear in Sections 4.7–4.9.

In order to realize fine-grained microscopic simulation, several agent-based concepts and models have been suggested. Many works related to agent-based car-following models were motivated by the simplicity and elegance of cellular automata-based models of traffic flow, such as the model by Nagel and Schreckenberg (1992). This cellular automata can be easily coined as an agent-based model describing the decision making of a simulated driver with only a few rules depending on the gap to the preceding vehicle and some additional random delay. Also, continuous car following characterized as driving strategy models can be coined as very simple agent-based modelling (see also Section 3.2). In these, a vehicle–driver unit adapts its speed depending on the distance and speed to the preceding entity.

The work by Burmeister *et al.* (1997) can be seen as one of the first agent architectures for traffic simulation, despite the fact that it tackles mostly the reactive driving level in form of an

---

[1] `http://www.citiesinmotion.com/`

agent-based car-following behaviour. A survey on driving models can be found in Kesting *et al.* (2009). Suggestions for driver models reproducing individual driving styles (with respect to individual parameters like preferred speed, acceptable gap to the preceding vehicle, acceleration rate, reaction times) can be found in Ehlert and Rothkrantz (2001). The vision in their work was to create more realistic traffic scenarios by including more heterogeneity in the model of the drivers.

A similar simulation goal can be found in Paruchuri *et al.* (2002), which describes a traffic simulation with different types of drivers (aggressive, cautious) in a situation where drivers do not need to adhere to traffic rules (chaotic intersection).

Other traffic simulators focus on simulating movement behaviour through intersections. For example, Doniec *et al.* (2008) rely on behaviour-based rules for the agents moving through intersections where actions that violate norms may result in traffic accidents. Espié and Auberlet (2007) show that standard car-following models lack anticipation beyond the immediately preceding entity for showing realistic behaviour especially in situations with ramp lanes. Their simulation model ARCHISIM allows complex agent-centric perception for realistic driving models; recent versions focus on realistic environmental representation involving also the perception of non-standard situations with emergency vehicles or unexpected obstacles (Ksontini *et al.*, 2012). Traffic accidents and improvement of pedestrian safety also motivate the work of Waizman *et al.* (2012) who use a 3D model with detailed agent manoeuvre behaviour and analysis of driver or pedestrian agents perception. Luo and Boloni (2012) focus on detailed perception models for simulating driving strategies with lane changes. Benenson and Martens (2008) require a realistic traffic flow simulation for the task of reproducing the search behaviour of drivers for free parking places in the residential areas of Tel Aviv. Agents move through the network and may or may not find free parking places.

There is a variety of agent-based microscopic traffic flow simulators, dealing with different levels of behaviours, in different configurations and granularity of traffic networks. It seems that their number has increased a lot during the last decade, thus is not possible to discuss each of them. We expect an increasing number of such traffic simulators, with more and more attractive visualization features, as the complexity involved in their implementation is decreasing. A recent trend—mainly for coping with scalability in large-scale simulations—are multi-level models such as suggested in Navarro *et al.* (2011) for urban (pedestrian) simulation, or attempts to use Graphics Processing Unit (GPU) computing (Strippgen & Nagel, 2009).

3.6 From traffic simulators to testbeds

Testing or benchmarking novel management and control approaches of traffic signals, toll policies, V2X communication, etc. are big drivers of traffic flow simulation. Only few traffic flow simulators exist that allow manipulation of different components in a flexible way as needed by research. Commercial as well as non-commercial microscopic flow simulation tools (see Barceló, 2010) for a compendium about the currently most important simulators)—such as VISSEM[2], PARAMICS[3], AIMSUN[4] or SUMO[5]—provide only restricted access to the actual controllers and simulation infrastructure. Barceló *et al.* (2004) discuss which elements are necessary in a traffic simulation testbed, whereas Timóteo *et al.* (2012) describe the TraSMAPI infrastructure for providing an uniform access to a traffic simulator.

4 Agent-based control and management

The increase in transportation demand can no longer be met by just providing additional capacity; this is economically, environmentally, or socially difficult to justify. Thus, the importance of

---

[2] `http://www.ptv-vision.com`

[3] `http://www.sias.com/ng/sparamicshome/sparamicshome.htm`

[4] `http://www.aimsun.com`

[5] `http://sumo.sourceforge.net/`

improving the existing infrastructure by means of a better utilization of the available capacity is growing. Most of the optimization methods target traffic control.

Several strategies of traffic control exist; almost all fit the control loop described by Papageorgiou (2003). The basic elements of this control loop are: (i) the physical network, its model, the model of demand, and disturbances (that can be measured, detected, or forecast); (ii) control devices (traffic signals, variable message signs, etc.); (iii) surveillance devices (e.g. loop detectors); and (iv) the control strategy. Some of these are computational entities: models, surveillance, and control strategies. While the former (models) has been discussed in Section 3, the next subsections focus on control strategies.

We start with some references for classical control of traffic signals (Section 4.1), included here in order to direct non-familiar readers to non-agent-based works, and to make clear which are their strengths and drawbacks. We have opted to cover mainly progressive systems[6]. Sections 4.2–4.8 then discuss several agent-based approaches and frameworks for traffic control. In Sections 4.2 and 4.3, the focus is on optimization techniques, from various perspectives (single intersection, networked with and without coordination among the various agents). Works discussed in these sections treat demand as an abstract entity; no driver or route choice is explicitly considered. In Section 4.4, we present and discuss learning-based approach to single intersections as well as networks of controllers, while Section 4.5 deals with approaches that consider self-organization in progressive systems. Section 4.6 deals with the literature that considers the demand by means of models of individual drivers that act and interact with the control infrastructure. In this section, we cover both optimization and learning-based approaches; in the latter case some deal with co-learning as both drivers and the controllers are represented by learning agents. Intersections not controlled by traffic signals are discussed in Section 4.7. Then, Sections 4.8–4.9 address mainly new technologies, approaches, and frameworks that aim at dealing with lightless intersections, and autonomous driving. Traffic management is then addressed in Section 4.10.

Before continuing, we remark that there is an expressive number of publications using various techniques from artificial intelligence such as genetic algorithms and fuzzy inference, which are not included here because they do not tackle the problem from a distributed and/or decentralized point of view. Thus, the criterion to select the works presented next was primarily whether or not techniques from multiagent systems were used or at least considered.

4.1 Classical algorithms and software

Optimization of traffic signal controllers in an off-line way is the basis of well-established algorithms such as TRANSYT (Robertson, 1969), which generates optimal coordinated plans for fixed-time operation. Of course the main drawback of this method is that plans are computed for a static situation, based on historical data. SCOOT—Split Cycle and Offset Optimization Technique (Hunt *et al.*, 1981)—is similar to TRANSYT but is traffic-responsive (uses data from buried detectors). SCATS—Sydney Coordinated Adaptive Traffic System (Lowrie, 1982)—is also based on real-time data. The main difference to SCOOT is that it is a hierarchical system. Similar frameworks that are, however, decentralized are Prodyn (Henry *et al.*, 1983), OPAC (Gartner, 1983), and UTOPIA (Di Taranto, 1989). TUC—*Traffic-responsive Urban Traffic Control* (Diakaki *et al.*, 2002)—models traffic flow using a variant of the store-and-forward control model to compute control policies in an efficient way.

The strengths of TRANSYT, SCOOT, and SCATS are their reliability, robustness, and large number of deployed systems. However, they also have some drawbacks. The simplest methodology for creating a signal-timing makes several assumptions, including the existence of a default value for saturation flow (normally one that reflects 'typical' conditions). However, this seldom

---

[6] Due to the fact that coordination has a broader meaning in multiagent systems, we use the terms progression or synchronisation to designate what in traffic engineering is denominated coordinated, synchronised, or progressive systems (popularly known as green waves).

happens in real-world cases. Besides this and the issue of cost and maintenance, there are several other reasons why those approaches may fail. In traffic networks without well-defined traffic flow patterns, for example, morning and afternoon peaks, that approach may not be effective. This is clearly the case in big cities where the business centers are no longer located exclusively downtown. Also, in some cities, 'secondary' streets turned out to be as important as traditional arterials due to saturation of main arterials and/or to the use of navigation systems (which may consider secondary streets in the suggested route). Traffic patterns can also be affected by accidents, floods, snow, etc. Finally, *a priori* determination of the appropriate signal plans for the different times of the day is a complex task that requires a lot of domain knowledge.

These issues show that simple off-line or even on-line optimization of the synchronization in *arterials* alone cannot cope with changing traffic patterns. This happens because traffic is a highly dynamic process, thus the optimal signal plan can hardly be determined in advance. Given the increasing volume of traffic, more flexible and robust approaches are necessary. In the next sections, we deal with agent-based traffic control and management approaches, which differ on their degrees of sophistication as well as on their scopes.

4.2 Optimization of single traffic controllers

In the work by Kosonen (2003), various agents interact in a single intersection using fuzzy inference. Each phase is modelled as an agent that can change the lights of the group to green when required by the volume of traffic and when permitted by other agents. Hence, there is a need for negotiation between agents about how to operate together. Although the focus seems to be on intra-controller coordination, Kosonen mentions that agents can also negotiate with neighbouring intersections, both upstream and downstream, by means of exchange of local traffic and control data. The paper does not provide more details about how the fuzzy control is performed when it comes to negotiate with neighbouring agents. Thus, the challenge here is how to integrate fuzzy control with coordination of agents, especially when it comes to solve both the local and the coordinated optimization problems.

Also Koźlak *et al.* (2008) have proposed a multiagent system to control traffic signals based on intersection managers (IMs; called intersection coordinators). Besides these managers, agents are also in charge of predicting future traffic conditions.

4.3 Optimization of network of controllers

In the next subsections, a more complex family of optimization problems is addressed, namely optimization of networks of agents located at neighbouring intersections. This departures from optimization of single controller and poses an extra problem because agents can either ignore other agents (and hence incur in sub-optimal solutions, at least from the point of view of the system-level performance), or coordinate with neighbour agents. Coordination is a key issue in multiagent systems and many propositions exist that can lead to improvement of performance in networks of controllers. It seems that a favourite technique is the use of communication to exchange data about traffic states and control actions. However, there is no obvious protocol that prescribes how the non-local information should be integrated with the local one. For instance what should be done when conflicts arise? This is the case if, for example, there is a request from a neighbour and this request is not compatible with the local actions an agent is trying to implement to solve its own local optimization problem.

Besides communication-based approaches, other popular ones are: (i) organizing control agents in a hierarchical structure so that conflicts are solved at an upper level; (ii) letting agents learn how to control their local environments or coordinate their actions; (iii) let agents self-organize. The first still needs communication at some extents (though less than a pure communication-based approach). The other two approaches, which appear in Sections 4.4 and 4.5, respectively, present the potential drawback that the time taken to learn or self-organize may be incompatible with the dynamics of the environment. In the particular case of learning-based approaches, unless agents receive utilities for their actions that are aligned with the global objectives of the system, the performance at global level tends to be poor.

4.3.1 Progressive systems

Classical approaches to progression or synchronization of traffic signals work mostly off-line and focus on synchronization of traffic signals in an arterial. The main difficulty to extend the synchronization to a network or to more directions of traffic is the fact that in some key intersections, conflicts may appear because different directions compete for bandwidth. The conventional approach is to let a traffic expert solve these conflicts. Alternatives to such approach seek to replace the traditional arterial green wave by shorter green waves in segments of the network. This can be done, for example, using negotiation over the question of which traffic direction shall be given more bandwidth. Although this requires communication among the agents, other approaches depart from this assumption (see Section 4.5).

An approach based on distributed constraint optimization problems (DCOP) was proposed in Oliveira *et al.* (2005). It is intended to be a compromise between that totally autonomous coordination with implicit communication that is described in Section 4.5 and the classical centralized solution. The constraints in this problem arise from the fact that, in each node of the graph, a traffic signal cannot coordinate to establish a progression with neighbours located in a different direction at the same time. A conflict occurs when two neighbours want to coordinate in two different traffic directions.

Junges and Bazzan (2008) have later extended this scenario to bigger networks, aiming at investigating computational issues related to performance, such as time to reach an agreement and number of exchanged messages.

4.3.2 Optimization of split

In order to improve the solution provided by the TUC approach (Section 4.1), the model predictive control approach is used by de Oliveira and Camponogara (2010) in order to handle, in conjunction with TUC, the constraints of the split optimization problem. Moreover, the authors also formulate the problem using techniques related to multiagent system to model the interconnection of sub-systems with local input constraints. Given dynamic equations and algebraic constraints, their framework formulates the optimization problem arising from the model predictive control and proceeds to decompose it into a network of coupled, small sub-problems to be solved by agents located at intersections. These agents sense the state variables and set the values of the control variables, communicating with agents in the neighbourhood in order to acquire data from these and coordinate their actions. Simulation studies corroborated the convergent behaviour of the solution proposed by the network of agents.

Apart from using pure communication to coordinate multiple agents, an alternative is to use some kind of organizational structure where an hierarchy of authority exists. Roozemond (2001) has proposed an agent-based urban signal control system that can be considered hierarchical and consists of several intersection control agents, authority agents, and road segment agents. Authority agents supervise and control several intersection agents, each of these in charge of managing the control at one intersection. Neighbouring intersection agents also communicate. Coordination is based on authority. Intersection agents use predetermined rules to select one control strategy and a prediction model that estimates the traffic situation in the near future. Based on these predictions and on a set of rules, the intersection agent calculates signal plans, check with adjoining agents, and plans the signal control strategy. This is sufficient to achieve good results in non-conflicting cases. The conflicting ones are passed to the authority agents. The author mentions that often an intersection agent should sacrifice some performance in favour of cooperative behaviour, and that some sort of negotiation is necessary. Also an advice-seeking mechanism is suggested but not detailed. Given that conflicts must be solved in real-time, this is likely to be a key. Besides, although the author claims that decisions are not centrally controlled, and that agents independently select the option with the highest (local) efficiency, this is valid only in non-conflicting cases. These cases may need the interference of the authority agent.

Another hierarchical multiagent system (here with three levels) is proposed by France and Ghorbani (2003). In the first level, local traffic agents (LTAs) represent intersections; these are responsible for providing appropriate traffic signal timing. Because a local optimum may not be so when observed from another perspective, there is a second level in the system, in which a coordinator traffic agent (CTA) supervises a few LTAs. According to the authors, a CTA should provide means by which the optimal local light pattern can be slightly modified to accommodate the global performance. This is hardly a trivial question. The paper only mentions that, by means of communication between the local components and the CTAs, and by means of storing all relevant information in a third level (controlled by an information traffic agent), it is possible to handle congestion caused by diverse natures. Another issue is that the paper does not mention what is a global solution in this case. If it is global in the sense of the CTA controlled area, then a situation may arise in which two or more CTAs compute their 'locally global' solution and, again, these may not be optimal from the network (thus, global) point of view.

In summary, split optimization has received a lot of attention by the multiagent systems community. While the single intersection case is relatively straightforward, big challenges are associated with the case involving several controllers because there is no obvious way to coordinate their selection of control actions. This can be performed independently, implicitly, or by means of communication. Especially in the latter case, it is an open question how to handle conflicts, except if hierarchically superior decisions are made and lower level agents just implement them. This, however, has the effect of reducing the autonomy of intersection agents and may even compromise local performances. Thus, more flexible coordination mechanisms are still a necessity.

4.4 Learning-based approaches: the compromise between independent versus centralized learning

A popular approach related to control of traffic signals is to let one or more agents learn a policy for mapping states to actions by observing the environment and selecting actions. Thus, most of the proposed approaches deal with reinforcement learning where the control problem is modelled by means of a Markov decision process (MDP). Depending on the actual formulation of the MDP, it may not be computationally feasible to solve the problem because the space of state-action pairs grows exponentially. This is due either to the discretization regarding the number of states in a single intersection, or to the number of intersections, or both. Solutions for this problem have been proposed in two directions. First, instead of modelling the whole problem of n intersections in a single MDP, one may use n independent learners. Second, it is possible to use function approximation to deal with the size of the space state. The former is a popular approach because it has the advantage of scaling to a large number of agents. However, because agents learn independently in an environment where other agents are doing the same, this approach may lead to global inefficiency or at least sub-optimality.

Next some of the proposed approaches are reviewed; we start with those based on a local, independent perspective of single agents, and later move to approaches based on multiagent learning or at least on some kind of interaction or coordination between the involved agents.

One of the first works dealing with reinforcement learning for traffic control was based on stochastic games (Camponogara & Kraus, 2003). The emphasis is on applying one of the following policies: uniformly random policy (assigns the same probability to all actions available to an agent); best-effort policy (lets traffic flow from the lane with the longest queue); and one achieved by using Q-learning. For the experiments, two intersection agents were used. Depending on the traffic conditions, if both agents implement Q-learning, there is a significant reduction in the waiting time.

In Silva *et al.* (2006), authors propose the use of model-based instead of model-free approaches to deal with the highly dynamic and non-stationary nature of traffic flow patterns. In this approach, it is assumed that flow patterns are non-stationary but they can be nearly divided in stationary dynamics that need not to be known *a priori*. Each model is assigned to an optimal policy (which is a mapping from traffic patterns to signal plans), and to an estimation of the quality of a given partial model. The creation of new models is controlled by a continuous evaluation of the prediction errors generated by each partial model. This model contains an estimated transition function and also an estimated reward function. If the environment changes and a local policy turns suboptimal (congestion increases over a threshold), then the system creates a new model. Whenever possible, the system reuses existing models instead of creating new ones. Although this mechanism was tested in a network of nine traffic signals, it remains an independent learning method and thus an extension is necessary to tackle joint states and joint actions.

In Prashanth and Bhatnagar (2011), the learning problem is solved in a centralized way, that is, a single MDP models all involved controllers. Thus, it is restricted to a grid of 3 × 3 controllers. Clearly, the centralized MDP poses a limitation for real-world applications. On the other hand, the authors focus on an important problem, namely the high dimension of the space state, by using a function approximation technique.

Given the issue of computation complexity of centralized MDP versus sub-optimality of independent learning, a recent trend is to let controllers learn independently but allow them to interact and combine their policies. This way one may depart from total centralization, as well as from total independence, and find a compromise between these two extremes. We now discuss some of these tentatives. These are mostly based on agents being able to coordinate their actions by means of several techniques such as communication of advice, collaboration, supervised learning, etc.

The question of how heterogeneous groups of agents can benefit from communication to improve their learning skills was investigated by Nunes and Oliveira (2004) in which information from several sources during learning is used in a simplified traffic control problem. Members of a team of agents may communicate among them or with members of other teams that are solving similar problems in different areas. Different types of agents use various techniques such as neural networks and heuristics.

Exchange of information by independent controllers is also tackled by Oliveira and Bazzan (2009) where controller agents may or may not benefit from exchanging information. Thus, this work analyses how agents can benefit from sharing information as well as the consequences of this cooperation in the performance of the traffic system. Authors show that having information is not sufficient if the other agents are not acting in a coordinated way; extra information can have a negative effect in the learning process if the information is unnecessary. Also, noise in the traffic patterns and the inter-dependencies among the agents' actions are relevant factors to consider when there is a need to decide between using shared information, using only local information, or an in-between solution.

The work of Kuyer *et al.* (2008) also focus on cooperative learning and explicit coordination among agents, that is, joint actions among them are considered. However, this leads to an increase in complexity. Thus, some simplifying assumptions must be made. In this case, authors assumed that the coordination takes place only among direct neighbours in the network. Under this assumption, the global coordination problem may be decomposed into a set of local coordination problems and can be solved with the use of coordination graphs as proposed by Guestrin *et al.* (2002). To find the optimal joint actions in such graphs authors apply an algorithm that estimates the optimal joint action. This work assumes that relationships among controllers are restricted to their direct neighbourhoods. However, in traffic this is not always the case. Although coordination graphs can be extended to other kinds of neighbourhoods, this approach has problems handling a high number of agents in each graph. Moreover, it is not clear what happens when controllers belong to more than one coordination graph as in this case the policies learned by the controllers in different graphs are conflicting.

In Bazzan *et al.* (2010), a layered architecture is proposed where supervisor agents are in charge of a small group of controller agents. This tries to balance the number of joint versus independent actions using a kind of organizational control in line with the work by Zhang *et al.* (2009).

4.5 Self-organizing progressive systems

As mentioned previously, one way to achieve progressive systems is to let agents self-organize regarding the decision about which green waves form. The issue of non-communication between traffic signals and local interaction was especially popular until some years ago when the cost of communication (e.g. wireless) was prohibitive.

In this sense, Oliveira *et al.* (2004) have proposed an approach based on swarm intelligence. Each intersection controller behaves like a social insect. Signal plans are seen as tasks to be performed by the insect without any centralized control or task allocation mechanism. Stimuli to perform a task or, sometimes, to change tasks, are provided by the vehicles that continuously produce 'pheromone' while waiting for green lights. Thus, the volume of traffic coming from one direction can be evaluated by the intersection agent, and this may trigger some signal plan switching. Experiments have shown that the agents achieve synchronization without any central management. However, the time needed to converge to a stable coordination can be high, which is a negative aspect especially in highly dynamic environments.

A similar idea appears in Helbing *et al.* (2005) where vehicles can synchronize traffic signals and organize green waves. Also in Gershenson (2007), traffic signals self-organize by means of three methods, also with no direct communication between them. It is shown that the adaptation to traffic conditions reduces waiting times and number of stopped vehicles.

Bazzan (2005) has used techniques of evolutionary game theory and stochastic games, in which there is no explicit communication among agents: traffic signal agents act in a dynamic environment having only local knowledge. However, these agents perform experimentation and receive a reward that depends also on the experimentation performed elsewhere in the neighbourhood. Stochastic events that may take place in the network are modelled by mutations. During the learning process, a fitness for each strategy is computed and influences the next generation of strategies that are used by the agents. Depending on the frequency of the stochastic events, agents are able to coordinate better towards the global goal. Results have shown that in scenarios where the volume of vehicles is nearly equal in both directions, the central progression does not perform well compared with the agent-based mechanism, because the latter is adaptive. A shortcoming of the approach is that payoff matrices (or at least the utilities and preferences of the agents) are required, that is, these figures have to be explicitly formalized by the designer of the system. This makes the approach time-consuming when many different options of coordination are possible. Thus, a challenge here is how to design mechanisms for large-scale of networks.

In Prothmann *et al.* (2011), an observer/controller architecture for signal control is proposed. This effort appears within the framework of organic computing (Schmeck, 2005). The idea is that a signal controller is the system under investigation, while an observer monitors the local traffic demand and evaluates the performance of the active signal plan. Based on the observations, the controller selects and optimizes signal plans by means of a two-levelled learning mechanism. In level 1, signal plans are selected on-line from a previous learnt mapping. This relies on a variant of a learning classifier system (XCS). Level 2 is in charge of optimizing the available signal plans when an unknown demand arises. This optimization is achieved by means of evolutionary algorithms. This approach was tested in isolated intersections as well as in progressive systems and with the inclusion of self-organized routing.

4.6 Approaches explicitly addressing demand

So far we have discussed traffic control approaches that focus on signal optimization and do not explicitly consider the demand. In the following, we discuss works aiming at the integrating both the demand and the control sides.

4.6.1 Co-learning based on waiting time

Some approaches for control at intersections propose that the waiting time of drivers is considered when computing or deciding the timing of traffic signals. In this line, Wiering (2000) describes the use of reinforcement learning by traffic signal agents in order to minimize the overall waiting time of vehicles in a small grid. Agents learn a value function that estimates the expected waiting times of vehicles given different settings of traffic signals. One interesting issue tacked in this research is that a kind of co-learning is considered: value functions are learned not only by the traffic signals, but also by the vehicles that can thus compute policies to select optimal routes to their respective destinations. This vehicle-based representation enables the estimations of reward values. The ideas and some of the results presented in this paper are important. However, it must be stressed that policies are computed by the traffic signal and not by the driver itself.

A similar RL-based method for controlling traffic signals is presented by Steingröver *et al.* (2005) to minimize the total travel time of all vehicles in the network. Thus, the control perspective is a global one, although the actions are local. Agents are the traffic signals but the learning task is formulated in a way that the state representation is vehicle based (waiting times for individual vehicles), aggregated over all vehicles around the intersection. As in Wiering (2000), it is assumed that the traffic signals are able to know the value functions of vehicles around (plus information such as destination). Another issue is that the more information about the individual vehicle, the bigger the state space. The paper also investigates other forms of state representation, with different learning abilities, generalization, and performances.

4.6.2 Co-evolution of drivers and controllers

In Bazzan *et al.* (2008), the learning task also involves two kinds of agents: drivers and traffic signals, each having its own goal and learning algorithm. The objective of local traffic control is obviously to minimize queues in a spatially limited area (e.g. around a traffic signal). The objective of drivers is to minimize their travel times. A typical commuting scenario modelled as a grid was used, where drivers repeatedly select a route to go from an origin to a destination. Thus, drivers have a large set of routes to select from. The control is done via decentralized traffic signals. The strategies of the traffic signals are: (i) keep the default signal plan; (ii) greedy (run green time for the phase or link with the higher occupancy); (iii) use single agent Q-learning. Regarding the driver, these can use three strategies: (i) select one of the possible routes randomly (each time it departs); (ii) select a route greedily (always pick the one with best average travel time so far); (iii) select a route in an adaptive way meaning that the average travel times so far are used to compute a probability to select the route.

Later, the authors have extended this method to incorporate on-the-fly re-routing (Bazzan & Klügl, 2008). Here drivers react to their perception of jammed links and adapt. This means that new routes are not necessarily learned. Although the authors have shown that re-routing may compensate an eventual inefficient traffic control (by the traffic signal agents), it remains an open question how this can be combined with reinforcement learning techniques by the drivers.

4.6.3 Learning for coordination of drivers' choices

Tumer and colleagues tackle congestion problems in a broad way using the metaphor of minority games. In Tumer *et al.* (2008), multiagent learning algorithms are applied to two formulations of the problem of selection of departure times so that they do not end up producing queues at certain times. These formulations are bottom-up and top-down. Different time slots have different desirability that reflect users' preferences. The system utility is measured from the perspective of a city manager that seeks to minimize system-wide delays. From the perspective of drivers, these aim at maximizing a personal objective function (e.g. the difference between the desired and the actual arrival time). In both cases, the fact that agents greedily pursue their best interest causes traffic to worsen for everyone. Agents' actions are determined based on a reinforcement learning algorithm. The key issue in Tumer *et al.* (2008) is to ensure that the agents receive utilities that promote good system-level behaviour. Authors emphasize that agents' utilities must be: aligned with the system utility (ensuring that when agents aim to maximize their own utility they also maximize the system's one); and sensitive to the actions of the agents so that these can determine the proper actions to select without having to explicitly consider all other agents in their particular utility functions.

In Tumer *et al.* (2009), a similar approach is used to explore the impacts of agent reward functions on two traffic problems: selection of departure time and selection of lane. The authors make an important remark about one issue that arises in traffic problems but does not arise in many other domains (e.g. rover coordination), namely ensuring that drivers follow the advice they receive. A related problem also arises when the city manager's reward is at odds with a social welfare function. Determining what incentives to provide to the agents so that these two seemingly different objectives can be simultaneously maximized is a critical problem that bears further study.

This interesting approach is perhaps more suitable to a scenario composed by autonomous vehicles because in this case rewards could be communicated. Human drivers do not have many means of perceiving reward, apart from, for instance travel time. However, travel time alone is not necessarily a good measure of global utility. Thus, a challenge here is to design appropriate mechanisms that ensure that drivers are getting utilities that are both aligned with the system reward, and are as sensitive as possible to changes in the reward of each agent.

4.7 Lightless

Mandiau *et al.* (2008) consider a single intersection, without traffic signals, having only stop signs as a two-player game played by autonomous vehicles approaching or entering the intersection. The actions available for the autonomous vehicles are to stop or to move; the types of games and corresponding payoffs are set depending on priority rules. The authors have performed tests for two agents, three agents, and with complex intersections (with a high number of conflicting movements). Results have shown that their model is able to reproduce the actual flow of vehicles. Limiting issues here are how to design such mechanisms to deal with large-scale networks.

The seminal work by Dresner and Stone (2004) has proposed a reservation-based intersection control where autonomous vehicles try to cross an intersection without conventional traffic signals. In that publication, it was assumed that: autonomous vehicles are not allowed to turn, do not change lanes, and all begin travelling roughly at the same speed. The reservation is performed as follows. First, each autonomous vehicle informs the IM: the time it will arrive at the intersection, the velocity, direction, maximum and minimum acceleration, and other vehicles properties. Then, the IM simulates the journey of the autonomous vehicle given the IM's knowledge about other similar reservations. If the space requested by the autonomous vehicle (for a given timeslot) is already occupied, then the request is rejected, in which case the autonomous vehicle must decelerate and try again. If the request is accepted, it must be kept or cancelled by the autonomous vehicle (in case it cannot be met).

Some low-level issues left open were: what happens if the driver has to make last second changes, whether the autonomous vehicle is really committed to the deceleration, what happens if conventional vehicles participate in the system as well, and what happens outside the range of the IM after autonomous vehicle's have to decelerate in response to a denial of request. In order to cope with these issues, in Dresner and Stone (2005, 2008), some of the previous assumptions were relaxed, especially the requirement that autonomous vehicles maintain a constant velocity in the intersection, and that they do not make turns. The improved protocol proposed is based on rules that the vehicles are expected to follow: a vehicle may not enter the intersection without a reservation; the vehicle must try to follow the actions prescribed by the IM; the autonomous vehicle cannot try to improve its own journey (as the manager will ignore a new request if the autonomous vehicle has already a reservation granted). Also, a modification in the protocol allows autonomous vehicles to share the system with conventional vehicles.

Finally, in Au *et al.* (2011), a further management policy is proposed, which puts the request messages on hold and then process several requests at once, in order to avoid starvation at secondary roads.

That seminal approach by Dresner and Stone (2004) plus its extensions have inspired other authors (see next section); major challenges are to address the low-level issues related to both fine grained activities of the driving task, and how to cope with mixed traffic (human and autonomous vehicles).

4.8 Market-based approaches

The reservation-based approach discussed in the previous section was also extended in Dresner and Stone (2006) to address deadlocks and delays caused by drivers putting reservations in a suboptimal way because the IM processes requests on a first come, first served basis.

This has also inspired the work of Vasirani and Ossowski (2009, 2011) where the reservation-based approach is extended to cover networks of intersections. The mechanism is market-based where drivers trade with the infrastructure agents in a virtual marketplace, purchasing reservations to cross intersections. Market rules were designed with the aim of aligning the global profit (revenues from the infrastructure use) with the social welfare (e.g. average travel time), in a way that, in situations of similar traffic load, an increase of the infrastructures monetary benefits usually implies a decrease of the drivers average travel times.

Schepperle and Böhm (2007, 2009) propose an approach that takes the valuation of the drivers into account. The focus is on single intersections. In Schepperle and Böhm (2007), a procedure is proposed, which consists of four steps: vehicle contacts the intersection; vehicle acquires an initial time slot to cross the intersection; if not satisfied, a vehicle can try to acquire a better time slot, this time from another vehicle; vehicles cross the intersection. In the second step, an auction (e.g. a second-price, sealed-bid) is run among the vehicles that do not yet possess a time slot. A variant is proposed in which a vehicle with a high valuation subsidize another vehicle that is located ahead and that is going to bid for a time slot but has a low valuation. Subsidizing vehicles ahead is expected to guarantee a better time slot for both vehicles. In the third step, vehicles arriving late can acquire time slots that have already been auctioned off. In Schepperle and Böhm (2009), the authors explicitly discuss some challenges for making valuation-aware control systems operational.

Balan and Luke (2006) have proposed history-based controllers intended to provide a kind of global fairness. According to the authors, the original inspiration was to allow drivers informing the traffic signals whether they are in a hurry or not. In the history-based approach, information about vehicles' immediate past histories is collected by the traffic signals, allowing them to communicate with other signals through the immediate histories of vehicles as these travel from intersection to intersection. They base their approach on the notion of historical fairness by allowing vehicles to store credits they receive when waiting at red lights, and cash the credits in when passing through intersections. Traffic signals base their decisions on the credits of various vehicles at the intersection. When a vehicle reaches its destination, it reports its average waiting time over all intersections and this is one metric used to assess the efficiency of the control. One issue is that vehicles have no interest in doing this and, if they do, there is no incentive to report the actual time taken.

4.9 Collaborative driving and cooperative car navigation

Collaborative driving is the focus of DesJardins *et al.* (2009). The authors describe an agent-based cooperative architecture that aims at controlling and coordinating vehicles, also showing that reinforcement learning can be used for this purpose. A multi-layered architecture was presented, which relies on both an action layer and on a coordination layer: the former is used to manage low-level vehicle control actions such as braking, accelerating, or steering, while the latter is responsible for high-level action choice by integrating cooperative decision making between vehicles. In both layers, multiagent reinforcement learning techniques are used, showing that the integration of these techniques at all levels of the autonomous driving controller yields efficient results for vehicle control and coordination. This approach thus facilitates the efforts of the system's designer, as the complex details related to vehicle control and related to the numerous possibilities of inter-vehicle interactions are automatically handled by the learning algorithm.

As already mentioned in Section 3.4.2, Yamashita and Kurumatani (2009) have proposed a cooperative car navigation system with route information sharing. Each vehicle transmits route information (current position, destination, and route to the destination) to a route information server and this estimates future traffic congestion using current congestion information. Estimates are then fed back to each vehicle, which uses them to re-plan its route. To evaluate the effect of the route information sharing system, two indexes are used: individual incentive and social acceptability. Results of the experiments have confirmed that a cooperative car navigation system with the route information sharing generally satisfied individual incentive and social acceptability, as well as had an effect on the improvement of traffic efficiency.

4.10 Agent-based management

Traffic management is another domain where there are opportunities for cooperating agents. This is probably due to the fact that agent-related technologies facilitate modularity, modelling, and simulation of management strategies. In this line, many tools, environments, and prototypes have been proposed. Apart from those already discussed in Section 3 (modelling and simulation), next we present some tentatives of developing integrated environments for traffic simulation and control, in order to better manage the traffic system as a whole.

In van Katwijk and van Koningsbruggen (2002), concepts related to agents and multiagent systems were used to formulate examples of traffic management based especially on heterogeneous agents and system components. From the modelling perspective, BDI was suggested (see also Section 3.4) to represent and reason about agents' knowledge. Also, for cooperative agents, a coordination protocol based on priority over resource allocation is discussed that helps the agents to distribute the flow of vehicles. Resources that agents can manipulate are mainly the flow of vehicles and allowed speed. Authors mention some kinds of resource-bounded decisions agents must make: amount of resource to request, time of the request, which priority the request has, contingency for non-granted requests and/or changes in priority order. The authors discuss the problems that are ultimately tied to conflict resolution and the problem of local versus global performance. They try to solve the latter by formulating a hierarchical organization where authority both solves conflicts and pursue global performance. There are, however, other solutions (see Section 4.3.2); in any case, these issues were tackled in the recent literature but remain partially unsolved. Thus, an obvious challenge here is how to solve conflicts that appear due to insufficient resources and about which model of organization to adopt, namely a hierarchical, authority-based one versus a totally decentralized with lesser commitment with global performance.

In van Katwijk *et al.* (2005), a testbed for traffic management is presented, which intends to manage different levels of complexity, a diversity of policy goals, and different forms of traffic problems. Their testbed aims at rapid development of multiagent-based management and control, consisting of an interaction model, intelligence models, and a world model. The former aims at modelling interactions among the agents, mainly via communication (see for instance, problems that fit this testbed in Section 4.3.2). Intelligence models are supported by rule-based inference, and by Bayesian inference (to model uncertainty agents may have). Authors claim that these three components are useful to implement decentralized traffic concepts that were reported in the literature. However, a closer analysis seems to indicate that their work is particularly suitable for hierarchical organized controllers and for inter-controller coordination. Therefore an evident challenge is to extend this testbed to accommodate other kinds of organizations of controllers, intra-controller coordination (e.g. among phases or groups of phases that exist in an intersection), as well as other forms of coordination such as those not explicitly based on communication (see Section 4.3.1).

In Wang (2008), an approach is presented as a generalization of the feedback control mechanism in control theory. Wang's approach (the TransWorld) is based on a connection between the actual transportation system and its artificial counterpart. This connection is manifold, that is, has many modes. In the learning and training mode, both the artificial and the actual systems are only loosely connected and the former serves as data center for learning operational procedures. In the experimentation and evaluation mode, the artificial system serves as a platform to conduct experiments and eventually predict the behaviour of the actual system. In the control and management mode, which is the most challenging one, both systems must be tightly connected (for instance, real-time connected). It is expected that the artificial system is able to replicate the actual behaviour. Eventual differences can be used to generate feedback control. The author emphasizes the need for parallel execution of both systems, and that the artificial one is seen as a generalization of adaptive control, whereas the AI-based artificial system replaces an analytic reference model.

The underlying idea of artificial transportation systems is also present in the MAS-T^{2}er Lab (Rossetti *et al.*, 2008) project, which is based on agent programming that can deal with the social and behavioural models, such as those discussed in Section 3.4.

Regarding the management of systems that include or are based on autonomous vehicles, in Bazzan *et al.* (2012), an agent- and market-based approach is proposed to manage a fleet of automated guided personal rapid transit vehicles (e.g. pods). This study considers some variants for both the processing of demands (trips) and for the routing. These variants are centralized versus decentralized, with or without en-route re-planning, and action based versus first in, first served basis. Results show, at least in some cases, the feasibility of a centralized service. However, communication as well as reliability and fault-tolerance are of course important issues.

5 Achievements and future challenges

In Sections 3 and 4, we have described applications of agent-based technologies to several areas of traffic and transportation engineering—mainly for reproducing real-world traffic in simulation and controlling the traffic flow.

Regarding modelling and simulation in general (Section 3), the different approaches presented aim at one main objective: realistically reproducing intelligent human behaviour and decision making in scenarios that may consider high-level tasks (e.g. route choice), as well as low level ones, as for instance, actual driving. In the former case, more sophisticated models may include issues such as daily activities (shopping, drop and pick kids at kindergarten, etc.). For this, it is necessary to include a high level of detail and flexibility into the models to be simulated. These details not only refer to the intelligence of the drivers, but also to an explicit treatment of complex environment: vehicles, control, and management of the infrastructure.

On the other side, agent-based approaches for control and management have turned out to be attractive in situations where a flexible and robust synchronization between different autonomously acting entities is required. The vision is to design distributed controllers that adapt themselves to the current traffic status in a timely fashion and thus optimize or learn about the overall traffic flow for a more efficient use of the existing infrastructure.

In Section 4, suggestions for intelligent infrastructures were presented, which allow autonomous driving, aiming at improving the throughput and safety. Section 4.10 has shown that the advantage of agents can be mainly seen in their responsiveness, robustness, and the potential for integrated services. Here, a better level of service for traffic participants as well as traffic managers is intended.

As it could be seen in this review, there are many attractive and promising ideas. However, most of the described approaches—in modelling and simulations, as well as in control and management—lie still on the conceptual model level or are quite restricted prototypes. Thus, real-world cases pose the main challenge with their size, complexity, and issues related to robustness to potential failures. From this picture, two concrete kinds of challenges can be derived: scalability and validation in their broadest sense. A third challenge comes from new application areas for agents in traffic and transportation that emerge due to the development of new technologies such as autonomous driving, vehicle to vehicle or vehicle to infrastructure communication.

Scalability is one of the major challenges that has to be addressed for deployment and application in realistic scenarios, even if computing power is increasing. The number of agents or the size of the network that can be feasibly tackled seems at first sight to be a software engineering or even only an implementation problem. These are quite generic problems that can be found in most large-scale simulations and are often addressed with distributed computing architectures. However, in agent-based traffic simulation large scale cannot be dissociated from the following problems. First, regarding error propagation, the following are sensible issues: will a slightly wrong acceptable gap on one link of the traffic network lead to a wrong agent decision making about speed on a neighbouring link and, as a consequence, to an invalid routing decision down the road? How to organize and restrict interactions between agents so that the model is robust and does not exhibit chaotic behaviour? What if not only driver agents interact, but also intelligent infrastructure or vehicles interact with each other? How can such a model be reasonably tested or explored?

Another aspect of scalability of an agent-based traffic simulation refers to the feasible integration of multi-level decision making: a simulated driver must deal with unconscious processes as well as intentional decision making. Finding realistic high-level decision-making models beyond pure rationality, that are suitable for traffic scenarios would be a challenge on its own. Its integration into a larger driver model architecture with multiple time scales is even more challenging. One can see that there is currently no framework that is both completely flexible and multiscale[7].

None of the two-layer agent architectures discussed in Section 3.4.3 has been applied in realistic, large-scale settings. The vision is then to reach an integrated, multi-level treatment of a driver model capable of: (i) operate the vehicle (possibly by means of intelligent support features); (ii) decision making about initiation of lane change; (iii) informed, en-route choice, embedded into an on-line location choice model based on flexible and realistic activity planning.

Clearly, not all these levels are necessary for all applications, but actually a multi-level treatment of drivers decision making would be essential in a number of traditional simulation studies. We give two examples. First, if routing is not taken into account, it makes little sense to make predictions about traffic jams derived from low-level traffic flow simulations, especially when dealing with non-trivial networks. Second, simulations for estimation of volume of pollutants and emissions should take into account the selection of transportation modes, which depends on multi-objective reasoning of CO_2-balance aware traffic participants, etc.

Considering more and more complex models, an interesting question is whether or not new technologies will make 'obsolete' the necessity of this complexity in agent-based simulations. For instance, consider that there is ubiquitous up-to-date information about the current traffic state, or that every vehicle-driver unit possesses perfect navigation capabilities (e.g. by using navigation devices with up-to-date traffic information), or that autonomous driving leads to compliant behaviour (e.g. with traffic rules as well as with route recommendation). In these cases, the basic assumptions for non-agent-based methods for simulating human choices (such as route and mode choice) would be fulfilled. However, as long as there is still a human-in-the-loop, then agent-based simulation is the only simulation paradigm able to capture issues such as human deciding not to follow route guidance or traffic rules.

Validation is a major challenge in agent-based simulation in general and this is not different when it comes to traffic systems. Without providing credible simulation models or sufficiently tested control and management systems, the innovative concepts will not be deployed in a real environment.

For modelling and simulation, validation refers to a thorough testing related to how good the model corresponds to the real world. Such a validation can only be done if sufficient and appropriate data are available. There are still unsolved problems in validation in traffic simulation, as it can be seen in the recent establishment of a EU-COST Action to this topic (COST Action TU0903 'Methods and tools for supporting the use, calibration and validation of traffic simulation models').

Agent-based traffic simulations poses special problems to validation as they involve informal findings from psychologists that can hardly be validated on a purely empirical basis. The available data are virtually the same as in the case classical simulation paradigms are employed. However, this data has to support a different class of assumptions, parameter settings, etc. Thus, agent-based traffic simulation models are inherently under-determined. The advantage of improved visualization capabilities can only partially compensate that. Additional sources for data supporting validation of agent-based traffic simulations have to be found. Currently, the most discussed 'new' data source is (individual-level) trace data from GPS systems or mobile devices.

In addition to this coarse summary of major challenges, the reader finds other particular issues that were discussed in Section 4. These can be summarized around the problem of the alignment of

---

[7] MATSim - introduced in Section 3—is definitely a huge step into that direction, but after construction and full refinement of a daily plan, the agent is committed to this plan so that it has to execute it without interruption. Thus, in fact, one cannot yet affirm that this approach allows agents to be completely flexible.

utilities, at both local and global levels, which ultimately relates to mechanism design. Many approaches address the problem about how to relate the global optimum to the individual agents objectives for driving the autonomous agents decision making towards the system-level optimum. However, none of the suggested concepts and systems gives sufficient answers to the important question that refers to how to balance the autonomy of the local decision-makers against the top-down directive. This is in fact the reason why vehicular traffic is quite different from traffic in computer networks: vehicles are driven by human beings with a high level of *autonomy*; packages can be routed while drivers are free to even ignore detour information aimed at diverting them from traffic congestion.

6 Conclusion

The more societies and economies become complex and coupled, the more difficult is the modelling and control of traffic. Traffic and transportation engineering now requires solutions that are related to information technology and control engineering. From the former, an important contribution is being given by technologies related to autonomous agents and multiagent systems. These technologies are becoming more popular in various applications related to traffic in the context of smart cities. This paper has presented a review on various works that address the cross-fertilization between the above-mentioned areas.

It has a different nature of previous works. First, in Teodorovic (2008), applications are presented that are related to swarm intelligence, such as ant colony optimization (ACO) and particle swarm optimization (PSO). One may think that an insect or a particle is an agent, though there are several questions related to autonomy and especially to levels of centralization required by both ACO and PSO (e.g. pheromone matrices that are normally centrally computed or common knowledge) whose discussion is beyond this review. Second, compared with a recent review on agent technology for traffic and transportation (Chen & Cheng, 2010), here our focus goes beyond technological issues. Rather, we also give sufficient insights to the underlying concepts.

Regarding the works we have reviewed, we have grouped them in two main categories: modelling and simulation, and control and management. In these categories, we focus on vehicular traffic, leaving pedestrian and crowd simulation, public transportation, air traffic control, maritime traffic simulation, and all topics around logistics for future work. We have decided to do this in order to keep the focus on an already extensive area. We have tried to provide a description of underlying concepts, paradigms, and methods. After, we went on analysing the achievements obtained so far, as well as the challenges ahead. These challenges concern not only subareas of computer science, information technology, and engineering, but also ITS-related areas as well. In order to tackle most of those challenges, it seems obvious that researchers from different areas of computer science, computer engineering, traffic and transportation engineering have to get together and join efforts in multidisciplinary teams.

Acknowledgments

Ana Bazzan is partially supported by CNPq and Franziska Klügl by VINNOVA in the VINMER program.

References

Arentze, T. & Timmermans, H. 2005. Representing mental maps and cognitive learning in micro-simulation models of activity-travel choice dynamics. *Transportation* **32**, 321–340.

Arentze, T. & Timmermans, H. 2008. Albatross: overview of the model, application and experiences. In *Innovation in Travel Modeling 2008 Conference*, Portland.

Au, T.-C., Shahidi, N. & Stone, P. 2011. Enforcing liveness in autonomous traffic management. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence*, Burgard, W. & Roth, D. (eds). The AAAI Press, 1317–1322.

Auld, J. & Mohammadian, A. 2012. Activity planning process in the agent-based dynamic activity planning and travel scheduling (ADAPTS) model. *Transportation Research Part A* **46**, 1386–1403.

Balan, G. & Luke, S. 2006. History-based traffic control. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, Nakashima, H., Wellman, M. P., Weiss, G. & Stone, P. (eds). ACM Press, 616–621.

Balmer, M., Cetin, N., Nagel, K. & Raney, B. 2004. Towards truly agent-based traffic and mobility simulations. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems, AAMAS*, Jennings, N., Sierra, C., Sonenberg, L. & Tambe, M. (eds). IEEE Computer Society, vol. 1, 60–67.

Balmer, M., Rieser, M., Meister, K., Charypar, D., Lefebre, N. & Nagel, K. 2009. MATSim-T: architecture and simulation times. In *Multi-Agent Systems for Traffic and Transportation Engineering*, Bazzan A. L. & Klügl, F. (eds). IGI Global, 57–78.

Barceló, J. (ed.) 2010. *Fundamentals of Traffic Simulation*. Springer.

Barceló, J., Codina, E., Casas, J., Ferber, J. L. & García, D. 2004. Microscopic traffic simulation: a tool for the design, analysis and evaluation of intelligent transport systems. *Journal of Intelligent and Robotic Systems* **41**, 173–203.

Bazzan, A. L. C. 2005. A distributed approach for coordination of traffic signal agents. *Autonomous Agents and Multiagent Systems* **10**(1), 131–164.

Bazzan, A. L. C., de Oliveira, D. & da Silva, B. C. 2010. Learning in groups of traffic signals. *Engineering Applications of Artificial Intelligence* **23**, 560–568.

Bazzan, A. L. C., de Oliveira, D., Klügl, F. & Nagel, K. 2008. Adapt or not to adapt—consequences of adapting driver and traffic light agents. In *Adaptive Agents and Multi-Agent Systems III*, Tuyls, K., Nowe, A., Guessoum, Z. & Kudenko, D. (eds). Lecture Notes in Artificial Intelligence **4865**, 1–14. Springer-Verlag.

Bazzan, A. L. C., do Amarante, M. B. & da Costa, F. B. 2012. Management of demand and routing in autonomous personal transportation. *Journal of Intelligent Transportation Systems* **16**(1), 1–11.

Bazzan, A. L. C. & Junges, R. 2006. Congestion tolls as utility alignment between agent and system optimum. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, Nakashima, H., Wellman, M. P., Weiss, G. & Stone, P. (eds). ACM Press, 126–128.

Bazzan, A. L. C. & Klügl, F. 2005. Case studies on the Braess paradox: simulating route recommendation and learning in abstract and microscopic models. *Transportation Research C* **13**(4), 299–319.

Bazzan, A. L. C. & Klügl, F. 2008. Re-routing agents in an abstract traffic scenario. In *Advances in Artificial Intelligence*, Zaverucha G. & da Costa, A. L. (eds). Lecture Notes in Artificial Intelligence **5249**, 63–72. Springer-Verlag.

Bazzan, A. L. C., Wahle, J. & Klügl, F. 1999. Agents in traffic modelling—from reactive to social behavior. In *Advances in Artificial Intelligence*, Extended version appeared in *Proceedings of the UK Special Interest Group on Multi-Agent Systems (UKMAS)*, Lecture Notes in Artificial Intelligence **1701**, 303–306. Springer.

Benenson, I. & Martens, K. 2008. Geosimulation of parking in the city. In *Proceedings of the 5th Workshop on Agents in Traffic and Transportation*, at AAMAS 2008, May 13, Estoril, Portugal, Bazzan, A. L. C., Klügl, F. & Ossowski, S. (eds), 29–35.

Bhat, C. R. & Koppelman, F. S. 2003. Activity-based modeling of travel demand. In *Handbook of Transportation*, Hall, R. (ed.). Springer, 39–65.

Burmeister, B., Doormann, J. & Matylis, G. 1997. Agent-oriented traffic simulation. *Transactions of the Society for Computer Simulation* **14**(2), 79–86.

Camponogara, E. & Kraus, W. Jr. 2003. Distributed learning agents in urban traffic control. In *The Portuguese Conference on AI*, Moura-Pires, F. & Abreu, S. (eds). Lecture Notes in Computer Science **2902**, 324–335. Springer.

Charypar, D., Nagel, K. & Axhausen, K. W. 2007. An event-driven queue-based microsimulation of traffic flow. *Transportation Research Record* **2003**, 35–40.

Chen, B. & Cheng, H. H. 2010. A review of the applications of agent technology in traffic and transportation systems. *IEEE Transactions in Intelligent Transportation Systems* **11**(2), 485–497.

Chmura, T. & Pitz, T. 2007. An extended reinforcement algorithm for estimation of human behavior in congestion games. *Journal of Artificial Societies and Social Simulation* **10**(2).

Davidsson, P., Henesey, L., Ramstedt, L., Törnquist, J. & Wernstedt, F. 2005. An analysis of agent-based approaches to transport logistics. *Transportation Research C* **13**, 255–271.

de Oliveira, D., Ferreira, P. R. Jr & Bazzan, A. L. C. 2004. A swarm based approach for task allocation in dynamic agents organizations. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems, AAMAS*. IEEE Computer Society, vol. 3, 1252–1253.

de Oliveira, L. B. & Camponogara, E. 2010. Multi-agent model predictive control of signaling split in urban traffic networks. *Transportation Research Part C: Emerging Technologies* **18**(1), 120–139.

de Palma, A., Ben-Akiva, M., Brownstone, D., Holt, C., Magnac, T., McFadden, D., Moffatt, P., Picard, N., Train, K., Wakker, P. & Walker, J. 2008. Risk, uncertainty and discrete choice models. *Marketing Letters* **19**, 269–285.

Desjardins, C., Laumônier, J. & Chaib-draa, B. 2009. Learning agents for collaborative driving. In *Multi-Agent Systems for Traffic and Transportation*, Bazzan, A. L. C. & Klügl, F. (eds). IGI Global, 240–260.

Di Taranto, M. 1989. UTOPIA. In *Proceedings of the IFAC-IFIP-IFORS Conference on Control, Computers, Communication in Transportation*. International Federation of Automatic Control, 245–252.

Dia, H. 2002. An agent-based approach to modeling driver route choice behaviour under the influence of real-time information. *Transportation Research C* **10**(5–6), 331–349.

Diakaki, C., Papageorgiou, M. & Aboudolas, K. 2002. A multivariable regulator approach to traffic-responsive network-wide signal control. *Control Engineering Practice* **10**(2), 183–195.

do Amarante, M. d. B. & Bazzan, A. L. C. 2012. Agent-based simulation of mobility in real-world transportation networks: effects of acquiring information and replanning en-route. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2012)*, Valencia, Spain, 1351–1352.

Doniec, A., Mandiau, R., Piechowiak, S. & Espié, S. 2008. A behavioral multi-agent model for road traffic simulation. *Engineering Applications of Artificial Intelligence* **21**(8), 1443–1454.

Dresner, K. & Stone, P. 2004. Multiagent traffic management: a reservation-based intersection control mechanism. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, Jennings, N., Sierra, C., Sonenberg, L. & Tambe, M. (eds). IEEE Computer Society, 530–537.

Dresner, K. & Stone, P. 2005. Multiagent traffic management: an improved intersection control mechanism. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, Dignum, F., Dignum, V., Koenig, S., Kraus, S., Singh, M. P. & Wooldridge, M. (eds). ACM Press, 471–477.

Dresner, K. & Stone, P. 2006. Multiagent traffic management: Opportunities for multiagent learning. In *LAMAS 2005*, Tuyls, K., Hoen, P. J., Verbeeck, K. & Sen, S. (eds). Lecture Notes in Artificial Intelligence **3898**, 129–138. Springer Verlag.

Dresner, K. & Stone, P. 2008. A multiagent approach to autonomous intersection management. *Journal of Artificial Intelligence Research* **31**, 591–656.

Ehlert, P. A. M. & Rothkrantz, L. J. M. 2001. A reactive driving agent for microscopic traffic simulation. In *Proceedings of the 15th European Simulation Conference, Prag, 2001*, Kerckhoffs, E. J. H. & Snorek, M. (eds). SCS Publishing House, 943–949.

Epstein, J. M. 2007. *Generative Social Science: Studies in Agent-Based Computational Modeling (Princeton Studies in Complexity)*. Princeton University Press.

Espié, S. & Auberlet, J. M. 2007. ARCHISIM: a behavioural multi-actor traffic simulation model for the study of a traffic system including ITS aspects. *International Journal of ITS Research* **5**, 7–16.

Ettema, D., Arentze, T. & Timmermans, H. 2011. Social influences on household location, mobility and activity choice in integrated micro-simulation models. *Transportation Research Part A* **45**, 283–295.

France, J. & Ghorbani, A. A. 2003. A multiagent system for optimizing urban traffic. In *Proceedings of the IEEE/WIC International Conference on Intelligent Agent Technology*. IEEE Computer Society, 411–414.

Gartner, N. H. 1983. OPAC—a demand-responsive strategy for traffic signal control. *Transportation Research Record* **906**, 75–81.

Gershenson, C. 2007. *Design and Control of Self-organizing Systems*. PhD thesis, Vrije Universiteit Brussel.

Gilbert, N. 2007. *Agent-based Models (Quantitative Applications in the Social Sciences)*. Sage Publications.

Gipps, P. G. 1981. A behavioural car-following model for computer simulation. *Transportation Research Part B* **15**, 105–111.

Grether, D., Chen, Y., Rieser, M., Beuck, U. & Nagel, K. 2008. Emergent effects in multi-agent simulations of road pricing. In *48th Congress of the European Regional Science Association, August 2008*, Liverpool, UK.

Guestrin, C., Lagoudakis, M. G. & Parr, R. 2002. Coordinated reinforcement learning. In *Proceedings of the Nineteenth International Conference on Machine Learning (ICML)*, Sammut, C. & Hoffmann, A. G. (eds). Morgan Kaufmann, 227–234.

Hackney, J. & Marchal, F. 2011. A coupled multi-agent microsimulation of social interactions and transportation behavior. *Transportation Research A* **45**, 296–309.

Han, Q., Arentze, T., Timmermans, H., Janssens, D. & Wets, G. 2009. A multi-agent modeling approach to simulate dynamic activity-travel patterns. In *Multi-Agent Systems for Traffic and Transportation Engineering*, Bazzan, A. L. & Klügl, F. (eds). IGI Global, 36–56.

Han, Q., Arentze, T., Timmermans, H., Janssens, D. & Wets, G. 2011. The effects of social networks on choice set dynamics: results of numerical simulations using an agent-based approach. *Transportation Research A* **45**, 310–322.

Helbing, D., Lämmer, S. & Lebacque, P. 2005. Self-organized control of irregular or perturbed network traffic. In *Optimal Control and Dynamic Games*, Deissenberg, C. & Hartl, R. (eds). Springer, 239.

Henry, J., Farges, J. L. & Tuffal, J. 1983. The PRODYN real time traffic algorithm. In *Proceedings of the International Federation of Automatic Control (IFAC) Conference*, Isermann, R. (ed.). IFAC, 307–312.

Horni, A., Nagel, K. & Axhausen, K. W. 2011. High-resolution destination choice in agent-based demand models. Working Paper 682, Institute for Transport Planning and Systems, ETH Zürich.

Horni, A., Scott, D. M., Balmer, M. & Axhausen, K. W. 2009. Location choice modeling for shopping and leisure activities with matsim: combining micro-simulation and time geography. *Transportation Research Record* **2135**, 87–95.

Hunt, P. B., Robertson, D. I., Bretherton, R. D. & Winton, R. I. 1981. SCOOT—a traffic responsive method of coordinating signals. TRRL Lab. Report 1014, Transport and Road Research Laboratory.

Joubert, J. W., Fourie, P. J. & Axhausen, K. W. 2010. Large-scale agent-based combined traffic simulation of private cars and commercial vehicles. *Transportation Research Record* **2168**, 24–32.

Junges, R. & Bazzan, A. L. C. 2008. Evaluating the performance of DCOP algorithms in a real world, dynamic problem. In *Proceedings of the 7th International Joint Conference on Automatic Agents and Multiagent Systems*, Padgham, L., Parkes, D., Müller, J. & Parsons, S. (eds). IFAAMAS, 599–606.

Kesting, A., Treiber, M. & Helbing, D. 2009. Agents in traffic simulation. In *Agents, Simulations and Applications*, Uhrmacher, A. & Weyns, D. (eds). Taylor and Francis, 325–356.

Klügl, F. & Bazzan, A. L. C. 2004. Simulated route decision behaviour: simple heuristics and adaptation. In *Human Behaviour and Traffic Networks*, Selten, R. & Schreckenberg, M. (eds). Springer, 285–304.

Klügl, F. & Bazzan, A. L. C. 2012. *Agent-based Modeling and Simulation*. AI Magazine.

Klügl, F., Bazzan, A. L. C. & Wahle, J. 2003. Selection of information types based on personal utility—a testbed for traffic information markets. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Gini, M., Ishida, T., Castelfranchi, C. & Lewis Johnson, W. (eds). ACM Press, 377–384.

Klügl, F. & Rindsfüser, G. 2011. Agent-based route (and mode) choice simulation in real-world networks. In *IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, Lyon, France, 22–29.

Kosonen, I. 2003. Multi-agent fuzzy signal control based on real-time simulation. *Transportation Research C* **11**(5), 389–403.

Koźlak, J., Dobrowolski, G., Kisiel-Dorohinicki, M. & Nawarecki, E. 2008. Anti-crisis management of city traffic using agent-based approach. *Journal of Universal Computer Science* **14**(14), 2359–2380.

Ksontini, F., Espié, S., Guessoum, Z. & Mandiau, R. 2012. Traffic behavioral simulation in urban and suburban—representation of the drivers' environment. In *Advances on Practical Applications of Agents and Multi-Agent Systems*, Demazeau, Y., Müller, J. P., Rodríguez, J. M. C. & Pérez, J. B. (eds). Springer, 115–125.

Kuyer, L., Whiteson, S., Bakker, B. & Vlassis, N. A. 2008. Multiagent reinforcement learning for urban traffic control using coordination graphs. In *Proceedings of the ECML/PKDD*, Antwerp, Belgium, 656–671.

Lowrie, P. 1982. The Sydney coordinate adaptive traffic system—principles, methodology, algorithms. In *Proceedings of the International Conference on Road Traffic Signalling*, Sydney, Australia.

Luo, Y. & Boloni, L. 2012. Modeling the conscious behavior of drivers for multi-lane highway driving. In *Proceedings of the 7th Workshop on Agent in Traffic and Transportation*, Valencia.

Machado, A. M. & Bazzan, A. L. C. 2011. Self-adaptation in a network of social drivers: using random boolean networks. In *Proceedings of the Workshop on Organic Computing (OC 11)*, Müller-Schloer, C., Schmeck, H. & Ungerer, T. (eds). ACM, 33–40.

Mandiau, R., Champion, A., Auberlet, J.-M., Espié, S. & Kolski, C. 2008. Behaviour based on decision matrices for a coordination between agents in a urban traffic simulation. *Applied Intelligence* **28**(2), 121–138.

McBreen, J., Jensen, P. & Marchal, F. 2006. An agent-based simulation model of traffic congestion. In *Proceedings of the 4th Workshop on Agents in Traffic and Transportation*, Bazzan, A. L. C., Chaib-Draa, B., Klügl, F. & Ossowski, S. (eds), 43–49.

Nagel, K. & Schreckenberg, M. 1992. A cellular automaton model for freeway traffic. *Journal de Physique I* **2**, 2221–2229.

Navarro, L., Flacher, F. & Corruble, V. 2011. Dynamic level of detail for large scale agent-based urban simulations. In *Proceedings of 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, K., Yolum, P., Sonenberg, L. & Stone (eds). IFAAMAS, 701–708.

Nunes, L. & Oliveira, E. C. 2004. Learning from multiple sources. In *Proccedings of the 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems*, Jennings, N., Sierra, C., Sonenberg, L. & Tambe, M. (eds). IEEE Computer Society, vol. 3, 1106–1113.

de Oliveira, D. & Bazzan, A. L. C. 2009. Multiagent learning on traffic lights control: effects of using shared information. In *Multi-Agent Systems for Traffic and Transportation*, Bazzan, A. L. C. & Klügl, F. (eds). IGI Global, 307–321.

de Oliveira, D., Bazzan, A. L. C. & Lesser, V. 2005. Using cooperative mediation to coordinate traffic lights: a case study. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and*

Multi Agent Systems (AAMAS), Dignum, F., Dignum, V., Koenig, S., Kraus, S., Singh, M. P. & Wooldridge, M. (eds). IEEE Computer Society, 463–470.

Ortúzar, J. & Willumsen, L. G. 2001. *Modelling Transport*, 3rd edition. John Wiley & Sons.

Panwei, S. & Dia, H. 2006. A fuzzy neural approach to modeling behavioural rules in agent-based route choice simulations. In *Proceedings of the 4th Workshop on Agents in Traffic and Transportation, AAMAS 2006, May 9*, Bazzan, A. L. C., Chaib-Draa, B., Klügl, F. & Ossowski, S. (eds), 70–78.

Papageorgiou, M. 2003. Chapter 8: Traffic control. In *Handbook of Transportation Science*, Hall, R. W. (ed.). Kluwer Academic Publisher, 243–277.

Paruchuri, P., Pullalarevu, A. R. & Karlapalem, K. 2002. Multi agent simulation of unorganized traffic. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. ACM Press, vol. 1, 176–183.

Prashanth, L. & Bhatnagar, S. 2011. Reinforcement learning with function approximation for traffic signal control. *IEEE Transaction on Intelligent Transportation Systems* **12**(2), 412–421.

Prothmann, H., Tomforde, S., Branke, J., Hähner, J., Müller-Schloer, C. & Schmeck, H. 2011. Organic traffic control. In *Organic Computing A Paradigm Shift for Complex Systems*, Müller-Schloer, C., Schmeck, H. & Ungerer, T. (eds). Springer, 431–446.

Rindsfüser, G., Klügl, F. & Freudenstein, J. 2004. Multi-agent simulation for the generation of individual activity programs. In *Application of Agent Technology in Traffic and Transportation*, Klügl, F., Bazzan, A. L. C. & Ossowski, S. (eds). Birkhäuser, 165–180.

Rindt, C. R., Marca, J. E. & McNally, M. G. 2002. An agent-based activity microsimulator kernel using a negotiation metaphor. Technical Report, Institute of Transportation Studies, University of California.

Robertson 1969. TRANSYT: a traffic network study tool. Rep. LR 253, Road Res. Lab.

Ronald, N. A., Arentze, T. A. & Timmermans, H. J. P. 2011. The effects of different interaction protocols in agent-based simulation of social activities. *International Journal of Agent Technologies and Systems* **3**(2), 18–32.

Roozemond, D. A. 2001. Using intelligent agents for pro-active, real-time urban intersection control. *European Journal of Operational Research* **131**(2), 293–301.

Rossetti, R. J. F., Bordini, R. H., Bazzan, A. L. C., Bampi, S., Liu, R. & Van Vliet, D. 2002. Using BDI agents to improve driver modelling in a commuter scenario. *Transportation Research Part C: Emerging Technologies* **10**(5–6), 47–72.

Rossetti, R. J. F., Ferreira, P. A. F., Braga, R. A. M. & Oliveira, E. C. 2008. Towards an artificial traffic control system. In *Proceedings of the 11th International IEEE Conference on Intelligent Transportation Systems (ITSC 2008)*, Beijing, China, 14–19.

Schepperle, H. & Böhm, K. 2007. Agent-based traffic control using auctions. In *Proceedings of the CIA*, Klusch, M., Hindriks, K. V., Papazoglou, M. P. & Sterling, L. (eds). Springer, 119–133.

Schepperle, H. & Böhm, K. 2009. Valuation-aware traffic control—the notion and the issues. In *Multi-Agent Systems for Traffic and Transportation*, Bazzan, A. L. C. & Klügl, F. (eds). IGI Global, 218–239.

Schmeck, H. 2005. Organic computing—a new vision for distributed embedded systems. In *ISORC*, Lee, S. (ed.). IEEE Computer Society, 201–203.

Silva, B. C. d., Basso, E. W., Bazzan, A. L. C. & Engel, P. M. 2006. Dealing with non-stationary environments using context detection. In *Proceedings of the 23rd International Conference on Machine Learning ICML*, Cohen, W. W. & Moore, A. (eds). ACM Press, 217–224.

Steingröver, M., Schouten, R., Peelen, S., Nijhuis, E. & Bakker, B. 2005. Reinforcement learning of traffic light controllers adapting to traffic congestion. In *Proceedings of the Seventeenth Belgium-Netherlands Conference on Artificial Intelligence (BNAIC 2005)*, Verbeeck, K., Tuyls, K., Nowé, A., Manderick, B. & Kuijpers, B. (eds). Koninklijke Vlaamse Academie van Belie voor Wetenschappen en Kunsten, 216–223.

Strippgen, D. & Nagel, K. 2009. Multi-agent traffic simulation with cuda. In *International Conference on High Performance Computing & Simulation, HPCS '09*, Leipzig.

Sun, Z., Arentze, T. & Timmermans, H. 2012. A heterogeneous latent class model of activity rescheduling, route choice and information acquisition decision under multiple uncertain events. *Transportation Research Part C* **25**, 46–60.

Teodorovic, D. 2008. Swarm intelligence systems for transportation engineering: principles and applications. *Transportation Research Part C: Emerging Technologies* **16**(6), 651–667.

Timmermans, H. 2005, editor *Progress in Activity-based Analysis*. Elsevier.

Timóteo, I. J. P. M., Araújo, M. R., Rossetti, R. J. F. & Oliveira, E. C. 2012. Using TraSMAPI for the assessment of multi-agent traffic management solutions. *Progress in Artificial Intelligence* **1**, 157–164.

Tumer, K., Welch, Z. T. & Agogino, A. 2008. Aligning social welfare and agent preferences to alleviate traffic congestion. *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, Padgham, L., Parkes, D., Müller, J. & Parsons, S. (eds). IFAAMAS, 655–662.

Tumer, K., Welch, Z. T. & Agogino, A. 2009. Traffic congestion management as a learning agent coordination problem. In *Multi-Agent Systems for Traffic and Transportation*, Bazzan, A. L. C. & Klügl, F. (eds). IGI Global, 261–279.

van Katwijk, R. & van Koningsbruggen, P. 2002. Coordination of traffic management instruments using agent technology. *Transportation Research Part C: Emerging Technologies* **10**(5–6), 455–471.

van Katwijk, R., van Koningsbruggen, P., De Schutter, B. & Hellendoorn, J. 2005. A test bed for multi-agent control systems in road traffic management. In *Applications of Agent Technology in Traffic and Transportation* Whitestein Series in Software Agent Technologies and Autonomic Computing,, Klügl, F., Bazzan, A. L. C. & Ossowski, S. (eds). Birkhäuser, 113–131.

Vasirani, M. & Ossowski, S. 2009. A market-based approach to reservation-based urban road traffic management. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Decker, K., Sichman, J., Sierra, C. & Castelfranchi, C. (eds). IFAAMAS, 617–624.

Vasirani, M. & Ossowski, S. 2011. A computational market for distributed control of urban road traffic systems. *IEEE Transactions on Intelligent Transportation Systems* **12**(2), 313–321.

Wahle, J., Bazzan, A. L. C. & Klügl, F. 2002. The impact of real time information in a two route scenario using agent based simulation. *Transportation Research Part C: Emerging Technologies* **10**(5–6), 73–91.

Waizman, G., Shoval, S. & Benenson, I. 2012. Micro-simulation model for assessing the risk of car-pedestrian road accidents. In *Proceedings of the 7th Workshop on Agent in Traffic and Transportation*, Valencia.

Wang, F.-Y. 2008. Toward a revolution in transportation operations: AI for complex systems. *IEEE Intelligent Systems* **23**(6), 8–13.

Wardrop, J. G. 1952. Some theoretical aspects of road traffic research. *Proceedings of the Institute of Civil Engineers* **2**, 325–378.

Wiering, M. 2000. Multi-agent reinforcement learning for traffic light control. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000)*, Stanford, CA, USA, 1151–1158.

Yamashita, T., Izumi, K. & Kurumatani, K. 2004. Analysis of the effect of route information sharing on reduction of traffic congestion. In *Application of Agent Technology in Traffic and Transportation*, Klügl, F., Bazzan, A. L. C. & Ossowski, S. (eds). Birkhäuser, 99–112.

Yamashita, T. & Kurumatani, K. 2009. New approach to smooth traffic flow with route information sharing. In *Multi-Agent Systems for Traffic and Transportation*, Bazzan, A. L. C. & Klügl, F. (eds). IGI Global, 291–306.

Zhang, C., Abdallah, S. & Lesser, V. 2009. Integrating organizational control into multi-agent learning. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, ao Sichman, J. S., Decker, K. S., Sierra, C. & Castelfranchi, C. (eds), 757–764. Budapest, Hungary.

Zhang, L. & Levinson, D. 2004. Agent-based approach to travel demand modeling. *Journal of the Transportation Research Board* **1898**, 28–36.

Zhu, S., Levinson, D. & Zhang, L. 2008. Agent-based route choice with learning and exchange of information. In *Transportation Research Board Annual Meeting 2008 Paper Nr.08-2152*.