

# A REACTIVE DRIVING AGENT FOR MICROSCOPIC TRAFFIC SIMULATION

*Patrick A.M. Ehlert and Leon J.M. Rothkrantz*  
*Knowledge Based Systems Group*  
*Department of Information Technology and Systems*  
*Delft University of Technology*  
*Zuidplantsoen 4, 2628 BZ Delft, the Netherlands*

E-mail: [P.A.M.Ehlert@its.tudelft.nl](mailto:P.A.M.Ehlert@its.tudelft.nl), [L.J.M.Rothkrantz@cs.tudelft.nl](mailto:L.J.M.Rothkrantz@cs.tudelft.nl)

## KEYWORDS

reactive agents, microscopic traffic simulation, multi-agent systems, driving behaviour

## ABSTRACT

Computer traffic simulation is important for making new traffic-control strategies. Microscopic traffic simulators can model traffic flow in a realistic manner and are ideal for agent-based vehicle control. In this paper we describe a model of a reactive agent that is used to control a simulated vehicle. The agent is capable of tactical-level driving and has different driving styles. To ensure fast reaction times, the agent's driving task is divided in several competing and reactive behaviour rules. The agent is implemented and tested in a prototype traffic simulator. The simulator consists of an urban environment with multi-lane roads, intersections, traffic lights, and vehicles. Every vehicle is controlled by a separate driving agent and all agents have individual behaviour settings. Preliminary experiments have shown that the agents exhibit human-like behaviour ranging from slow and careful to fast and aggressive driving behaviour.

## INTRODUCTION

In the last two decades, traffic congestion has been a problem in many countries. To reduce congestion, most governments have invested in improving their infrastructure and are exploring new traffic-control strategies. A problem is that infrastructure improvements are very costly and each modification must be carefully evaluated for its impact on the traffic flow. Computer traffic simulations form a cost-effective method for making those evaluations. In addition, traffic simulations can evaluate the improvements not only under normal circumstances, but also in hypothetical situations that would be difficult to create in the real world. Obviously, the used simulation model needs to be accurate in modelling the circumstances and in predicting the results.

Intelligent agents can be used to simulate the driving behaviour of individual drivers. The adaptability and flexibility of intelligent agents allows them to control various types of vehicles with different driving styles. Each agent is equipped with its own behaviour settings to simulate personalised driving behaviour. This way, the

simulated vehicles will behave realistically and the interaction between multiple drivers can be studied.

This paper describes a model of a reactive agent that can perform tactical-level driving. Tactical-level driving consists of all driving manoeuvres that are selected to achieve short-term objectives. Based on the current situation and certain pre-determined goals, the agent continuously makes control decisions in order to keep its vehicle on the road and reach its desired destination safely.

## MICROSCOPIC TRAFFIC SIMULATORS

Many traffic simulators that are used today are macroscopic simulators. Macroscopic simulators use mathematical models that describe the flow of all vehicles. These models are often derived from fluid dynamics and treat every vehicle the same. Only the more advanced models can differentiate between vehicle types (e.g. cars, trucks, and busses) and even then all vehicles are treated equally within one vehicle type.

In real life many different types of vehicles are driven by different kind of people, each with their own driving style, thus making traffic flow rather unpredictable. In microscopic simulations, also called micro-simulations, each element is modelled separately, allowing it to interact locally with other elements. For example, every simulated vehicle can be seen as an individual with the resulting traffic flow being the emergent behaviour of the simulation. Microscopic traffic simulators are able to model the traffic flow more realistically than macroscopic simulators.

A Multi-Agent System (MAS) [Ferber 1999] can be used to form the basis of a microscopic traffic simulator. The main components (agents) of a multi-agent traffic simulator will be the vehicles. Every vehicle is controlled by an individual agent. Other important elements of the simulated environment can also be modelled as agents, for example a traffic-light agent that controls a group of traffic lights. In 1992 Frank Bomarius published a report on such a MAS [Bomarius 1992]. His idea was simply to model all the used objects as agents that could communicate the relevant data. Four years later two MSc students at the University of Edinburgh implemented this idea for their final MSc project [Chan 1996], [Chong 1996]. Their nameless text-based simulator uses Shoham's AGENT-0 architecture [Shoham 1993] to create multiple agents that function as vehicles or traffic lights, but also as roads and intersections. As the emphasis of their project was on creating a MAS-simulation and not necessarily creating realistic driving behaviour, all

their vehicle agents use very simple rules based on gap acceptance and speed. More advanced behaviours like overtaking cannot be modelled due to the simplicity of both their agent and simulation environment.

A more advanced simulation environment is the SHIVA simulator, which stands for Simulated Highways for Intelligent Vehicle Algorithms [Sukthankar et al 1996]. The SHIVA simulator was especially designed to test tactical-level driving algorithms and allows fast creation of different test scenarios. In his PhD thesis Rahul Sukthankar describes a reasoning system for tactical-level driving called POLYSAPIENT that was tested with SHIVA [Sukthankar 1997]. A drawback of the SHIVA simulator is that it needs a special SGI machine to run and is not publicly available.

At first glance, the approach we used with our driving agent resembles the POLYSAPIENT reasoning system used by Sukthankar, but its implementation is quite different. First of all, our simulator implements an urban environment. SHIVA and most other traffic simulators model highway or freeway traffic. Second, with our agent multiple behaviour parameters can be set to produce the desired driving behaviour. Most other simulators only use one or two driving-behaviour parameters (usually gap acceptance and preferred speed or an aggression factor) or none at all. Third, by using relatively independent behaviour rules our agent's functionality can be expanded or altered easily and the agent can be used in completely different environments.

## TRADITIONAL VERSUS REACTIVE AGENTS

Traditional intelligent agent architectures applied in artificial intelligence use sensor information to create a world model [Wittig 1992], [Rao and Georgeff 1995]. The world model is processed by common search-based techniques, and a plan is constructed for the agent to achieve its goal. The plan is then executed as a series of actions. This traditional approach has several drawbacks. Sensor constraints and uncertainties cause the world model to be incomplete or possibly even incorrect, and most traditional planning methods cannot function under noisy and uncertain conditions. Furthermore, in complex domains like tactical driving it is infeasible to plan a complete path from the initial state to the goal state, due to the large amount of searchable states and the inability to perfectly predict the outcome of all possible actions. The amount of possible states 'explodes' if realistic manoeuvres such as aborted lane changes and emergency braking are taken into account. As a result a real-time response cannot be guaranteed, making the traditional planning methods unsuitable for tactical driving.

Reactive agents, also called reflex or behaviour-based agents, are inspired by the research done in robotic control. Their primary inspiration sources are Rodney Brooks' subsumption architecture [Brooks 1986] and behaviour-based robotics [Arkin 1998]. Reactive agents use stimulus-response rules to react to the current state of the environment that is perceived through their sensors. Pure reactive agents have no representation or symbolic model of their environment and are incapable of foreseeing what is going to happen. The main advantage of reactive agents is

that they are robust and have a fast response time, but the fact that pure reactive agents do not have any memory is a severe limitation. This is the reason that most reactive agents use non-reactive enhancements.

## DRIVING AGENT MODEL

We have designed a model of a reactive driving agent that can control a simulated vehicle. The agent is designed to perform tactical-level driving and needs to decide in real-time what manoeuvres to perform in every situation. These decisions are based on the received input from the agent's sensors. After the agent reaches a decision, the instructions are translated into control operations that are sent to the vehicle.

The driving agent is modular in design. Every part can be adapted, replaced or otherwise improved without directly affecting other modules. The agent consists of: several sensors to perceive the environment, a communication module, a memory for storing data and controller for regulating access to the memory, a short-term planner, multiple behaviour rules and behaviour parameters, and an arbiter for selecting the best action proposed by the behaviour rules. A picture of the agent's layout is shown in Figure 1.

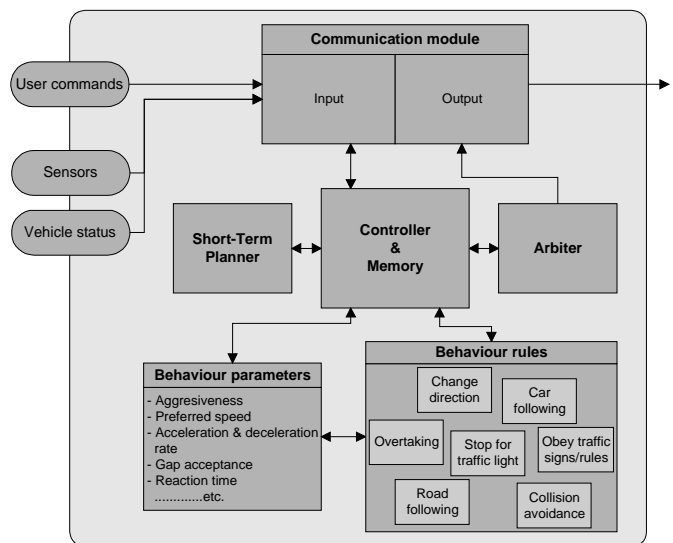


Figure 1: Driving agent layout

Our agent uses both traditional and reactive methods to perform its task, but the emphasis is on the latter since fast response times are important. Sensor information is stored in the memory and forms a temporary world model. Reactive procedures called behaviour rules or behaviours use the available information in the memory to quickly generate multiple proposals to perform a particular action. Planning in the traditional sense is not applied. The short-term planner only uses simple linear extrapolation to calculate the expected positions of moving objects and the arbiter determines the best available action based on the priority ratings of the action proposals included by the behaviour rules. We will discuss the agent's reasoning process, behaviour rules and behaviour parameters in more detail in the next subsections.

## Reasoning

The complete loop from receiving sensor messages to sending an output message to the vehicle can be seen as one reasoning cycle. The timing of a reasoning cycle and the activation of the agent's parts are done by the controller that also regulates the access to the memory. Since we want the driving agent to react in at least real-time, the agent is able to complete several reasoning cycles per second. The activation of the agent's parts is shown in Figure 2.




Figure 2: The reasoning cycle regulated by the agent's controller

The agent uses two types of sensor information. The first type gives information about the agent's environment, for example the distance and angle to objects, or the width of the agent's current lane. The second sensor type returns information about the agent's vehicle. This includes speed, acceleration, heading, wheel angle, and fuel level. In addition, the agent can receive orders from the user.

All information that is sent to the agent is received by the agent's communication module that contains knowledge of the used protocols. When a message is received, the communication module tries to recognise the message format, the sender and its content. When the message is ok, the input section of the communication module temporarily stores it until all received messages can be written to the agent's memory. Temporary storage is necessary since one does not want data in the memory to be read and written at the same time. Outgoing messages can be sent immediately since no conflicts can arise there. Next, all incoming messages are transferred to the agent's memory and the short-term planner makes a fast prediction of the position of all moving objects in the environment. Then the actual reasoning of the agent is performed by the behaviour rules, also called behaviours. They specify what to do in different situations. Based on the available data in the agent's memory, every behaviour can propose an action. All action proposals have a tally or priority rating. The arbiter selects the best proposal based on the priority ratings and sends it to the communication module. Finally, the communication module translates the proposal to control instructions that can be understood by the vehicle.

## Behaviour rules

The agent's driving task is divided into several subtasks that are automated by independent behaviour rules. This way the agent's functionality can be expanded easily without any modifications to the existing behaviours. The used behaviour rules are very much dependent of the agent's environment. Instead of a highway environment often used in traffic simulations, we have chosen to let the agent drive in an urban environment. The reason for this is

that an urban environment is one of the most difficult and complex traffic scenarios. In a city a lot of unexpected events can happen and the agent has to deal with many different situations. This way we can show the potential of our driving agent concept. Note that the design of our agent does allow driving in other environments. Only the agent's behaviour rules might need to be adapted or expanded. For our city environment we designed the following behaviours:

### *Road following*

The road-following behaviour is responsible for keeping the agent driving on the road. Besides controlling the lateral position of the agent's vehicle, based on the distance to the road and lane edges, the road-following behaviour also influences the agent's speed. It makes sure that it slows down for curves and on straight roads it will accelerate until the desired speed set in the agent's behaviour parameters is reached.

### *Intersection / changing directions*

If the agent approaches an intersection, its speed is reduced, precedence rules are applied, and the agent will choose one of the sideroads. Usually, this direction is chosen randomly, but it can also be set by the user. The changing-directions behaviour can be split up into several sub-behaviours, one for each type of intersection (e.g. intersections with or without traffic lights, or a roundabout). This is consistent with the fact that humans use different strategies to handle different types of intersections.

### *Traffic lights*

The traffic-lights behaviour makes sure that the agent stops for red or yellow traffic lights if possible. The behaviour checks if the sensed traffic light regulates the agent's current lane and slows down the vehicle. The agent's start-point for braking depends on its preferred braking pressure (deceleration rate) and is stored in the behaviour parameters.

### *Car following*

The car-following behaviour ensures that the agent does not bump into any other vehicle. If another car is driving in front of the agent, speed is reduced to match that car's speed. The precise braking pressure depends on the speed difference between the agent's vehicle and the other vehicle, the distance between them, and the set gap acceptance of the agent.

### *Overtaking and switching lanes*

Related to the car-following behaviour is the overtaking-and-switching-lanes behaviour. If a slower vehicle is in front of the agent, it may decide to overtake this vehicle. This decision depends on the velocity difference between the two vehicles and the available space to overtake the vehicle, both in front and to the left of the other vehicle.

### *Applying other traffic rules*

Besides traffic lights and precedence rules at junctions, other traffic rules need to be followed. Examples are, not driving at speeds above the local maximum, driving on the right side of the road as much as possible (in the Netherlands), and no turning in one-way streets.

For the traffic-rules behaviour it is necessary to keep track of the traffic signs and restrictions encountered by the agent. Because the memory of the agent will clear data on a regular basis to save space, the traffic-rules behaviour needs to keep track of these signs itself, in its own private memory space. This memory space is embedded within the behaviour. Note that the behaviour also needs to keep track when the signs and rules apply. Usually, turning onto a new road will reset most of the current restrictions.

#### Collision detection and emergency braking

The collision-detection and emergency-braking behaviour is a special kind of safety measure that is activated when the agent is on a collision course with an object. It can be seen as the human reflex to brake if something pops up unexpectedly in front of the vehicle. The behaviour tries to ensure that the vehicle can be halted at all times before it hits an object. Actions from the emergency-braking behaviour have the highest priority and always overrule all other behaviours.

#### Behaviour parameters

In order to create different driving styles all behaviour rules are influenced by behaviour parameters. One of the most important (visible) parameter is the driver's choice of speed. This choice has a large effect on the different driving subtasks. Drivers that prefer high speeds are more likely to overtake other vehicles than slower drivers and usually brake harder. Another implemented factor is the distance the driver keeps to other cars, also called gap acceptance. Aggressive drivers keep smaller gaps than less aggressive drivers. A third parameter is the driver's preferred rate of acceleration or deceleration. Again, aggressive drivers tend to accelerate faster than less aggressive drivers.

Besides the above-mentioned behaviour factors, other aspects can influence an agent's driving behaviour, for example the reaction time of an agent and the range of its sensors. An agent's reaction time can be altered by changing the length of its reasoning cycle. The sensor range determines the visibility of the agent and can be used to simulate fog or bad weather conditions.

## IMPLEMENTATION

We have constructed a prototype traffic simulator program to test our driving agent design. The programming language we used to build the simulator is Borland Delphi 5 Professional for NT. We have chosen this language in part since we were already familiar with it, but mainly because Delphi is an easy language, very suitable for quick prototyping.

Our simulator uses a kinematic motion model that deals with all aspects of motion apart from considerations of mass and force. The model implements smooth motion of vehicles, even during lane changes. Furthermore, the vehicles can move along realistic trajectories, but since forces are not modelled, the vehicles will perform manoeuvres without slipping.

#### The prototype simulator

The simulator program roughly consists of four elements: a user interface to provide visual feedback, a simulation controller, an environment containing simulated objects, and the driving agent model. The task of the simulation controller is to start, pause or stop a simulation run, and keep track of the elapsed time. The controller also initialises, starts and stops the used driving agents. During simulation, the controller regularly sends an 'update' order to the environment. The environment then calculates the new values for all its objects and sends relevant visual feedback to the screen. This 'simulation update' loop is shown in the left part of Figure 3. By default the update frequency is about 20 times per second, but this rate can be adjusted so that the program can be run on slower computers.

The environment is formed by all the simulated objects together. Different environments can be loaded via Map Data Files. These files contain a description of a road network and traffic control systems. Loading a Map Data File initialises the environment and data about the simulated objects described in the file is stored in the environment. Our current simulator implementation contains multi-lane roads, intersections, traffic lights, traffic light controllers and vehicles.

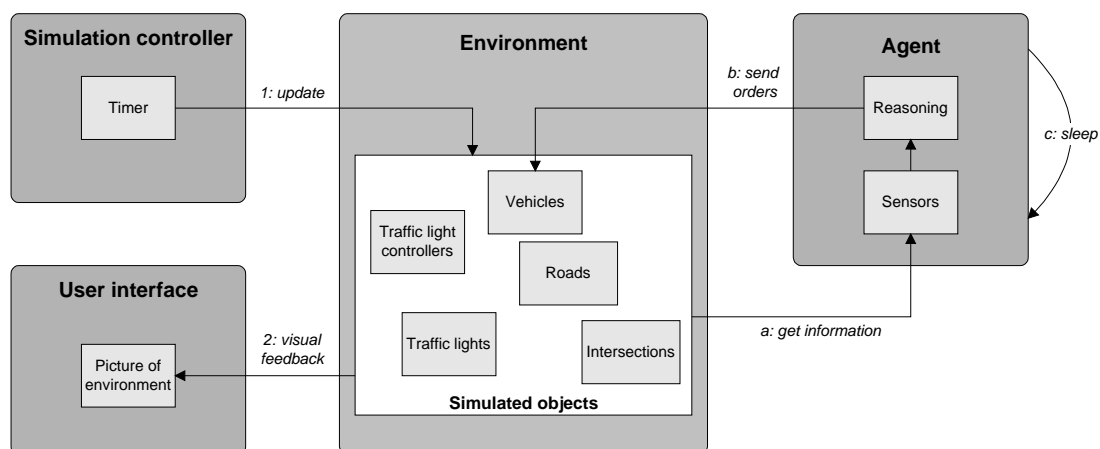


Figure 3: Simulation and agent update loops

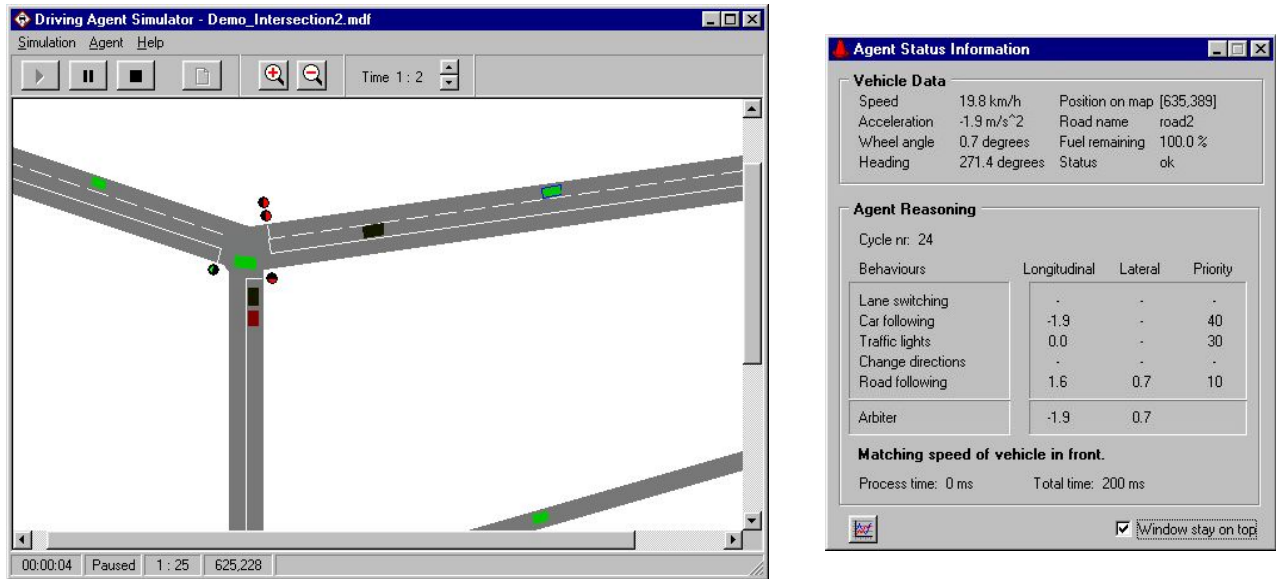


Figure 4: Screen shot of the prototype simulator used to test the driving agent

### The driving agent

Every vehicle in the environment has its own driving agent, but there is one agent that has the focus of attention and can be ‘controlled’ by the user. This means that the user can change the settings of this agent’s behaviour parameters and can follow its reasoning process in the Agent Status Information window shown in Figure 4.

All agents are implemented as threads, initialised by the simulation program. The advantage of using threads is that the simulation can be faster, running threads in parallel (if the operating system allows it), and that the agents can run independent of the simulation program. The disadvantage is that there is a limit to the number of threads one can use, because the overhead in managing multiple threads can impact the program’s performance. The execution loop of an agent is shown in the right part of Figure 3. After the agent finishes a reasoning cycle its thread is put asleep for the rest of its cycle time, which is set by the simulation controller. This is done to prevent agents from using all available CPU time. By default an agent’s cycle time is 200 ms, so the agents will perform 5 reasoning cycles per second.

The implementation of the behaviour rules is done using if-then rules. All behaviours are divided into several tasks. Tasks are executed in a serial manner, the least important task first and the most important task last. This way the important tasks ‘override’ the action proposals of less important tasks. The execution of the behaviour rules is also done consecutively, but in this case the execution order does not matter since the arbiter will wait until all behaviours are finished determining their action proposal.

### RESULTS AND DISCUSSION

We have presented a model of a reactive driving agent that can be used to control vehicles in a microscopic traffic simulator. A prototype simulation program was constructed to test our agent design. Although we have not validated the

used parameters yet, preliminary experiments have shown that the implemented agent exhibits human-like driving behaviour ranging from slow and careful to fast and aggressive driving behaviour. Here we present the results of one of our experiments, done using the first five behaviour rules discussed earlier in the “behaviour rules” section. The experiment consists of two different drivers approaching an intersection and stopping in front of a red traffic light. Both drivers perform this task without any other traffic present. The first driver is a careful driver with a low preferred speed, reasonably large gap acceptance and a low preferred rate of deceleration. We call this driver the ‘grandpa’ driver. The second driver is a young and aggressive driver, with a high preferred speed, small gap acceptance and a high preferred rate of deceleration. The drivers start at the same distance from the intersection. The speed of both vehicles during the experiment is shown in Figure 5.

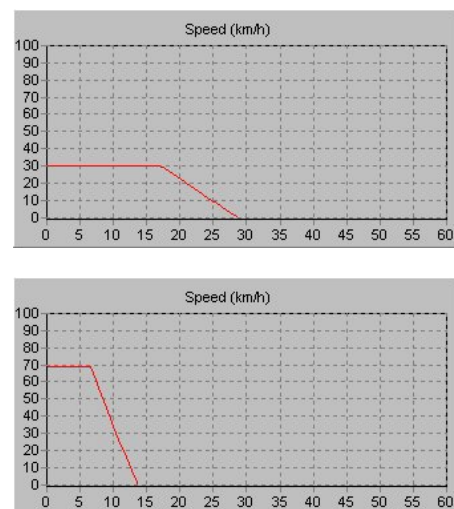


Figure 5: Speed of the grandpa driver (top) and young aggressive driver (bottom) during the experiment

Since the grandpa driver is driving at a lower speed, it takes a while before he starts braking, but his braking start-point (50m) is closer to the intersection than that of the young aggressive driver (65m), due to his lower speed. The difference between the used braking pressures is clearly visible. Both drivers brake with a relatively stable deceleration (approximately  $0.7 \text{ m/s}^2$  and  $2.7 \text{ m/s}^2$ ), which is consistent with human braking behaviour.

The experiment was done several times, but in almost all cases the shown graphs were roughly the same. Also, the precise stopping positions of both vehicles were approximately the same in all experiments. The young aggressive driver had a tendency to brake relatively late and often came to a stop just in front or on the stopping line. The grandpa driver on the other hand always came to a full stop well ahead of the stopping line. The stopping positions of both vehicles during one of the experiments is compared in Figure 6.




Figure 6: Compared stopping positions of the young aggressive driver (red vehicle) and grandpa driver (blue vehicle)

The aim of our simulation program was to test the design and functionality of our driving agent, but as a result its current implementation is rather inefficient since we did not optimise it for speed. Our main focus was on the correctness of the agent's driving behaviour and reasoning process. The computer used to implement and test the program is an Intel Pentium III, 450 MHz with 64 MB of RAM, running the Microsoft NT 4.00 operating system. On this computer we were able to run experiments with up to 30 vehicles. Experiments with more vehicles are possible, but result in a slow-running simulation. For this we are working on improving the simulator's memory management and processing speed.

A drawback of our simulator is that some unrealistic assumptions were made. Agent perception is perfect. All agents have a field of view of 360 degrees and objects are not obscured or covered by other objects. Further, vehicle actions are atomic. For example, braking is applied instantly after the action is sent to the vehicle. In real life this would occur more gradually. Also, pedestrians, cyclists and crosswalks are not yet modelled so the agent's ability to react to unexpected events was not yet accurately tested.

## CONCLUSIONS AND FUTURE WORK

The main advantage of agent-based microscopic traffic simulation over the more traditional macroscopic simulation is that it is more realistic. Instead of using general traffic-flow models, traffic becomes an emergent property of the interaction between agents. Another advantage is that agent-based simulation is more flexible. Changes to traffic scenarios can be made quickly by altering the position of individual vehicles and changing agent parameters. A disadvantage is the increase of computational resources and the higher number of parameters that need to be set and validated.

Preliminary experiments have shown that our driving agent exhibits human-like driving behaviour and is capable of modelling different driving styles, ranging from slow and careful to fast and aggressive driving behaviour.

At the moment we are experimenting with different types of agents in several scenarios. The goal is to study the current possibilities of our traffic simulator and agent in order to improve them further. The simulation environment should be made more realistic by adding new objects, such as busses, trucks, emergency vehicles, pedestrian crossings, cyclists, traffic signs, trees and buildings. Once the simulator is improved with the new objects the agent's functionality must be extended to deal with these objects. In addition, the simulation environment needs to be validated. Although we have tried to use realistic values for vehicle acceleration, turn radius, road size etc., the used settings might prove to be inaccurate. We also need to study human driving behaviour more extensively in order to validate our driving style models.

The drawback of the proposed improvements will be that both the simulation environment and the agent will need more computation time and will run more slowly. Therefore, we are considering using a distributed approach in the future so that the driving agents can run on different computers. The simulation controller and environment can act as a server and the agents can be the clients communicating to the server.

## REFERENCES

- Arkin, R. C. (1998) *Behavior-based robotics*. The MIT Press, Cambridge, Massachusetts.
- Bomarius, F. (1992) "A Multi-Agent Approach towards Modelling Urban Traffic Scenarios." Research Report RR-92-47, Deutsches Forschungszentrum für Künstliche Intelligenz, September 1992.
- Brooks, R.A. (1986) "A robust layered control system for a mobile robot." MIT AI Lab Memo 864, September 1985. Also in *IEEE Journal of Robotics and Automation* Vol. 2, No. 1, March 1986, pages 14-23.
- Chan, S. (1996) *Multi-agent Traffic Simulation – Vehicle*. MSc dissertation, Department of Artificial Intelligence, University of Edinburgh.
- Chong, K.W. (1996) *Multi-Agent Traffic Simulation - Street, Junction and Traffic light*. MSc dissertation, Department of Artificial Intelligence, University of Edinburgh.

- Ferber, J. (1999) *Multi-agent systems: an introduction to distributed artificial intelligence*. Addison Wesley Longman Inc., New York.
- Rao, A. S., and Georgeff, M. P. (1995) "BDI Agents: From Theory to Practice." In *Proceedings of the 1st International Conference on Multi-Agent Systems (ICMAS-95)*, San Francisco, USA, June 1995, pages 312-319.
- Shoham, Y. (1993) "Agent-oriented programming." In *Artificial Intelligence 60*, pages 51-92.
- Sukthakar, R., Hancock, J., Pomerleau, D. Thorpe, C. (1996) "A Simulation and Design System for Tactical Driving Algorithms". In *Proceedings of artificial intelligence, simulation and planning in high autonomy systems*.
- Sukthakar, R. (1997) "Situation awareness for tactical driving." Phd Thesis, Technical report CMU-RI-TR-97-08, Robotics institute, Carnegie Mellon University, January 1997.
- Wittig, T. (1992) *ARCHON: an architecture for multi-agent systems*. Ellis Horwood Limited, England.

## **AUTHOR BIOGRAPHY**

**PATRICK EHLERT** was born in Gouda, the Netherlands and has recently obtained his Master's degree in Computer Science at the Delft University of Technology. He just started working as a PhD student at the same institution.

**LEON ROTHKRANTZ** has a degree in psychology and mathematics and is working as a lecturer at the Delft University of Technology.