

Traffic Shifts in a Comprehensive Truck Toll System

A Real-World Multi-Agent Based Simulation Case Study

Dr. Heinz Schild
Bonn - Bad Godesberg, Germany

Summary

As part of a project which deals with comprehensive free-flow toll systems for trucks and buses, a multi-agent based simulation model was developed. This model is used here to discuss the general development process for multi-agent based simulations which map real-world situations and whose results are the foundation for real-world decisions. It will be demonstrated that developing such models is a demanding task, requiring detailed understanding of the relevant reality and meticulous work.

Goals

Based on a case study this paper aims to

- describe the approach when developing a multi-agent based simulation model of a real-world problem
- spell out the issues requiring special consideration in the course of the development and implementation process of the case study
- explain conclusions derived from the simulation results for real-world problems.

The case study is part of a project in which a concept for a comprehensive free-flow toll system for trucks and buses in central and eastern European countries is explored. In this concept ¹ all trucks and buses with a gross weight of 3.5 t or more are to be eligible for tolls on every public road of a country. Rebates are granted for kilometres driven on higher-ranking roads as an incentive to transfer traffic to these roads. Higher-ranking roads permit higher speeds, normally give rise to fewer accidents and have better noise protection facilities; their externality costs are normally below those of lower-ranking roads. Free-flow operations require that onboard units be installed in all the toll-eligible vehicles wanting to apply for a rebate on the standard toll charge².

¹ Schild, Heinz: Umfassendes Lkw-Mautkonzept, Begründung, Gestaltung und Beschaffung, Bonn, Oktober 2005, www.caselab.info

² Onboard units are needed to register the kilometres driven on the different road types. Two techniques are common nowadays. With the DSRC technique, data about the road segments which have been passed are collected with the help of gantries which communicate via microwaves with the onboard unit in the vehicle. With the GPS technique, the road segments passed are located using GPS (in future Galileo) and data are transmitted to the central system using GSM.

Only GPS-based onboard units are considered in this study. They give toll system operators greater freedom in determining the roads on which rebates will be granted.

Fig. 1 shows the main actors and processes of the relevant real-world segment as a UML Use-Case diagram ³.

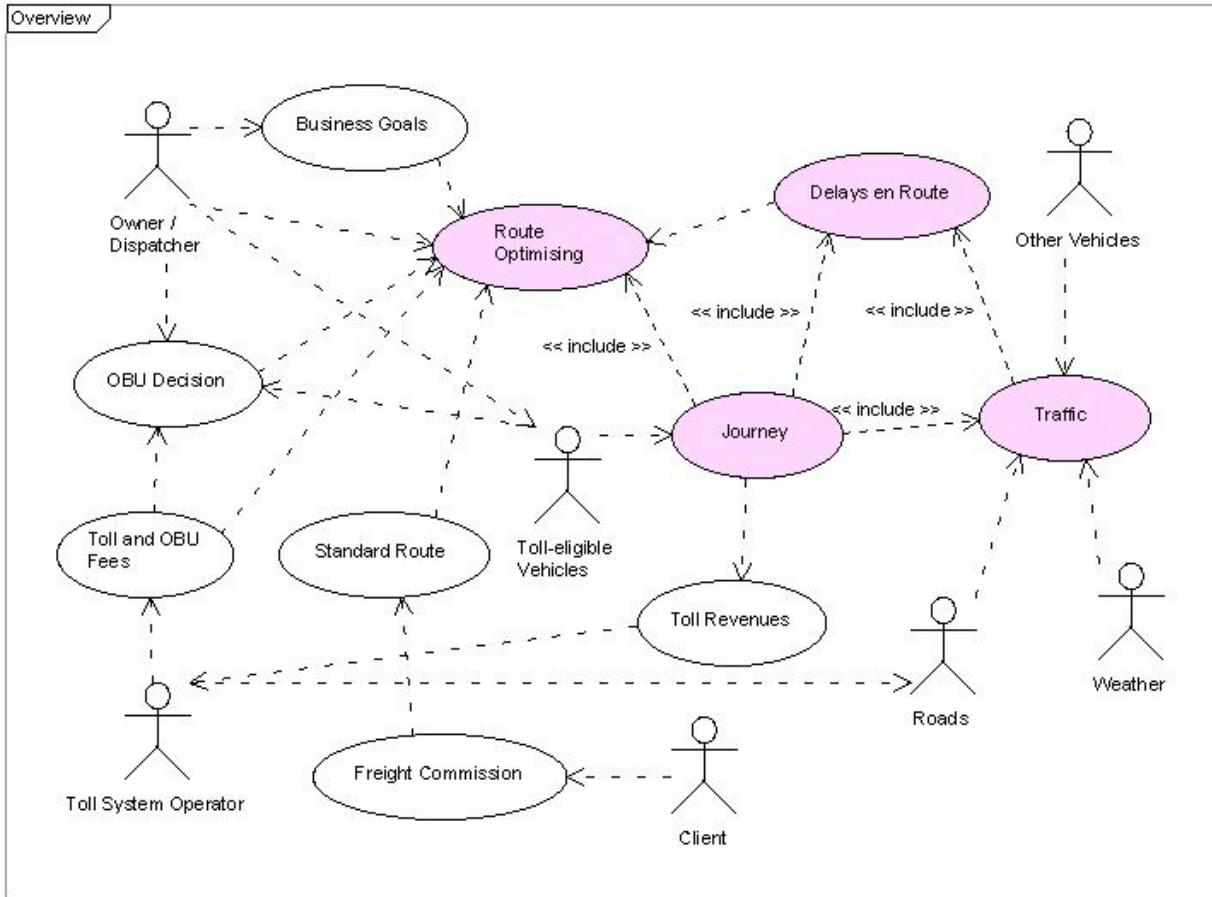


Fig. 1: Actors and processes

A client assigns a freight commission to the operator of a toll-eligible vehicle. The operator’s dispatcher plans the route in such a way that his business goals can be optimally fulfilled. Planning of the route depends on whether an onboard unit is installed in the vehicle or not. The operator will only install an onboard unit in a vehicle if he expects that during the lifetime of this vehicle the toll savings made by taking advantage of the rebates on higher-ranking roads will be greater than the costs for the onboard unit. The planned route may deviate from the standard route to take account of both the rebates on the standard toll charges and the expected delays due to congestion on the higher-ranking roads. Traffic conditions during the journey mainly depend on the amount of traffic on the roads used and their capacity. These traffic conditions may cause substantial delays which the dispatcher of the vehicle will bear in mind when he plans future routes for the vehicle.

³ UML stands for unified modeling language. Here version 2 was used as implemented in “Poseidon for UML Standard Edition from www.Gentleware.com, 2005

It is the intention of the case study to analyse and quantify the deviation of optimised routes from standard ones when traffic transfers from lower- to higher-ranking roads depend on the amount of rebate on the standard toll charges. The findings are important for toll system operators, as they can help them to estimate the amount of revenue they may lose due to rebates and to recognise the amount of traffic transferring to the higher-ranking roads which may subsequently generate political pressure to upgrade these roads.

It will be shown later that a multi-agent based simulation model seems to be the appropriate means to support this study.

Design and implementation

Characteristics and problems of multi-agent models

A multi-agent simulation model maps explicitly the behavior of every agent as influenced by its own attributes and the behaviour of other agents into a computer program which calculates the interactions of all agents over some period of time. Agents can be autonomous, proactive, reactive and interactive⁴. In our case agents are vehicles and road segments.

Multi-agent models sometimes have advantages over models which use global or aggregated entities. The agents, their attributes, behaviour and interactions may be directly observable in the real world. However, the modelling of their behaviour should be based on a relevant and broadly accepted behaviour theory and not on episodic and subjective impressions only. Models based on global or aggregated entities are highly abstract and depend on theoretical concepts and statistical data. The behaviour of their entities is not directly observable but deduced using statistical means, mostly regression analysis.

Nevertheless, the mapping of agents, their attributes and behaviour is normally not a straightforward task. Often constraints exist concerning the available development time, the expressive power of the simulation language and the know-how of the developer. Multi-agent simulations normally need much more computer power than global simulations and may be handicapped by the available computer systems. These constraints often do not allow all aspects of the real world to be mapped in their original scales. For instance in the real world modelled here, 40,000 toll-eligible vehicles drive on the roads simultaneously. For reasons discussed later, the model has to make do with 500 vehicles, meaning a scale of 1 : 80 has to be applied.

⁴ See also Wooldridge and Jennings (1995) quoted in Gilbert, Nigel, Troitzsch, Klaus G.: Simulation for the Social Scientist, Second Edition, Maidenhead, 2005, Chapter 8

In order to extract behaviour procedures directly from observations and pass them on into the model, scales must be chosen in such a way that their combination maps those states for which the behaviour is relevant on a 1 : 1 scale.

Even if the design of the scale follows this rule, mapping errors are unavoidable. They are caused by simplifications and by formulating behaviour relationships or selecting values for attributes which only approximate reality. Using a multi-step validation and calibration process, the effects of the mapping errors must be reduced enough to allow conclusions to be drawn from the model results which have relevance for the real world.

Graphical representations showing the agents and their states during a simulation run are very helpful in debugging a multi-agent simulation program and validating behaviour relationships. They are an essential part of a multi-agent simulation.

Overview of the design and implementation process

The design and implementation process of a multi-agent simulation model and program, though broadly similar to that of common simulation models, is in some respects quite different. It consists of a set of tasks which must be performed at least once⁵. The process is not always linear and may loop back.

The tasks are:

1. Describe the problem which is to be solved by a model in a comprehensive and clear way.
2. Select the appropriate modelling method and the computer tool for implementing it.

The following tasks are relevant if the multi-agent simulation method and a computer tool specially provided for this method are being used.

3. Identify the actors and use cases which cover the reality to be modelled and which appear relevant to the problem at hand. Select those processes which will build the core of the model, those actors which will become the agent classes and finally designate the processes and actors which affect the model via input parameters.
4. Determine the main output variables of the model and the main parameters of the alternative scenarios to be investigated with the model.
5. Specify the agent classes and subclasses of the model and their specific attributes.

⁵ See Gilbert, Nigel, Troitzsch, Klaus G.: Simulation for the Social Scientist, Second Edition, Maidenhead, 2005, Chapter 9

6. Determine the number of instances of the agent classes and the size of the simulation interval, bearing in mind the limits of the “simulation view window”⁶ which can be programmed with the chosen tool.
7. Design the possible movements of the dynamic agents in the simulation view window and their states.
8. Map the endogenous and interactive behaviour of the agents using programmable behaviour functions and applying approved behaviour theories, published observations of relevant specialists or own experiences and intelligent guesses.
9. Determine the time regime of the simulation in such a way that the output variables are not unduly disturbed by start-up abnormalities and oscillations.
10. Implement the simulation program and debug the code; search for mapping errors and implausibilities by studying the simulation view window and other ad-hoc charts and improve the program until its results no longer contradict relevant information about the real world.
11. Perform simulation runs by systematically varying parameter settings to explore the behaviour space of the model. Select one scenario whose results correspond reasonably with the real world. Fine-tune it by calibrating the behaviour equations of the agents until progress in the simulation and the final results sufficiently map the real world.
12. Perform simulation runs with the fine-tuned model and alternative scenarios to generate the behaviour space for the problem. The behaviour space should allow one to draw conclusions as to how, and how sensitively, the real world will react to exogenous decisions or events.
13. Document the model and its implementation and the evaluated scenarios with their inputs and results to substantiate the findings.

Step 1: Problem description

The problem which is to be dealt with in this case study has already been outlined above (see Fig. 1). It is always good practice to apply UML and its visualisation tools.

In this study the operator of a comprehensive truck toll system wants to examine how rebates on the standard toll charges which are granted for the kilometres driven on higher-ranking roads may affect the traffic flow and toll revenues. A computer model is to be developed to provide further insights.

⁶ In the “simulation view window” the agents and their moves during a simulation run are visualised. Agents are depicted as icons. Colours may be used to show different states of the agents.

A metric must be applied which is compatible with the real situation of a central European country in 2002. In this country there existed 2,000 km of motorways and 10,000 km of through roads. About 40,000 toll-eligible vehicles on average travelled on these roads during main traffic hours.

Step 2: Method and tools

The movement of vehicles and their interactions while on a journey will be simulated with the computer model. The traffic situations depend on the characteristics of the instances of the road classes and the different behaviour patterns of the vehicles. It seems obvious to apply the multi-agent based simulation method for the computer model and not one based on aggregated entities.

Several software packages are available for the implementation of multi-agent simulation models. Because of previous knowledge ⁷ and positive experiences with it ⁸, NetLogo for Mac OS X 10.4 on an Apple PowerBook G4 with a 1024 x 768 12" screen was selected. This constrains the maximum number of agents and the maximum number of simulation periods and simulation runs. The constraints must be considered when designing and implementing the model.

The statistics language R ⁹ was used for the preparation of some of the input data and the behaviour equations needed in the simulation model. The analysis of the results of the simulation runs was done with Quantrix Modeller ¹⁰, a multi-dimensional spreadsheet software.

Step 3: Core process and agent classes

The core process of the simulation model covers the four use cases which are shown in violet in Fig. 1. The main use case is "Journey". It includes "Route optimising", "Traffic" and "Delays en route". From among the actors directly associated with the core use cases, only those actors who interact in a way relevant to the problem form the agent classes. They are the "Toll-eligible vehicles", the "Other vehicles" and the "Roads".

The other actors and use cases affect the model through parameters which either are direct inputs to the simulation or are attributes of the agents or parameters of behaviour equations. Examples of the first type are the standard toll charges and the costs for onboard units. Other parameters are the specific mean journey length for every toll-eligible vehicle or the mean percentage of vehicles using a

⁷ Schild, Heinz: Enforcement in Free-Flow Systems: a Case Study, System Dynamics Society Proceedings of the 22nd International Conference, Oxford, 2004

⁸ Wilensky, U. (1999). NetLogo. <http://ccl.northwestern.edu/netlogo/>. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL

⁹ R : Copyright 2005, The R Foundation for Statistical Computing, Version 2.1.1 (2005-06-20), ISBN 3-900051-07-0, <http://cran.r-project.org/>

¹⁰ Quantrix Modeler, Professional Edition, Version 2.000.3, www.quantrix.com

specific instance of a road class. Further parameters are needed for the weights with which current and historical information or rumours about delays are combined to estimate the risk of a delay during the next journey.

Step 4: Result variables and alternative scenarios

During the design and implementation process of a multi-agent model one can easily be led astray. This can happen if one investigates too many interesting-looking details or concentrates primarily on those features which are easy to implement using a subset of the tool one is most familiar with. To prevent getting caught in a dead end, the variables for which the simulation must produce values should be defined before implementation work is started. These variables should be reporter variables, as they are called in NetLogo. Reporter variables do not have equivalents in reality but combine values from other variables for reporting purposes.

Here we are interested in the transfer of traffic from lower to higher-ranking roads as a result of rebates on the standard toll charges. Therefore the simulation model should deliver values for the following reporter variables:

- proportion of toll-eligible vehicles with onboard units installed
- mean relative saving of toll fees by taking advantage of rebates and transferring parts of the route to higher-ranking roads
- mean proportion of toll-eligible vehicles for which their dispatchers notice delays and take these into consideration when planning the next journey
- mean proportion of kilometres driven by toll-eligible vehicles which were transferred to motorways compared with all kilometres driven on motorways by these vehicles
- mean proportion of kilometres driven by toll-eligible vehicles which were transferred to through roads compared with all kilometres driven on through roads by these vehicles,
- mean proportion of toll-eligible vehicles driving on motorways.

In order to solve this problem, values of the reporter variables are needed for a set of alternative scenarios. The scenario parameters which define the behaviour space of the problem should be

- rebate schema (combination of the rebates for through roads and motorways)
- number of toll-eligible vehicles
- number of other vehicles.

Step 5: Agent classes and their instances

As already mentioned, three agent classes should be sufficient. They are

- the toll-eligible vehicles (trucks)
- the other vehicles (cars)
- the roads.

The next question is whether these classes should be split up into subclasses and if so, how. When answering this question, one should bear in mind that subclasses can be substituted by additional attributes of the instances. On the other hand, it is easier to visualise subclasses or their instances than attributes. Visualisation helps one to understand the behaviour of the model and makes debugging quicker but normally demands more computer resources.

Splitting up the class of toll-eligible vehicles into subclasses would be reasonable if the rebates depended on the weight or number of axles of the vehicles. But this is not an assumption in this study (although in reality it may be the case). Here different rebates are granted according to road type only. The instances of the class of toll-eligible vehicles are depicted in the simulation view window as small icons with a shape like a truck. They own static and dynamic attributes. For example, the attribute “mean journey length” is static because it does not change during a simulation run whereas the attribute “current speed” is dynamic. The attribute “OBU”, which may have the value “yes” or “no”, is set after the start-up period and then remains static. The icons of vehicles with an onboard unit are red, those without an onboard unit are blue. The distribution of the values of the static attributes of all instances should be sufficiently similar to that observable in the real world.

The role of the other vehicles in the model is to generate realistic congestion situations like those found during holiday periods. Therefore no subclasses are needed. The instances of the class of other vehicles are depicted as small yellow icons with a shape like a car. They also have static and dynamic attributes.

The class of roads must definitely be split up into subclasses. Following the practice in traffic engineering, three subclasses should be used ¹¹:

- motorways (and dual carriageways) or limited-access facilities (US)
- through roads (often called A roads) or arterials (US)
- local roads (all other roads including secondary urban roads) or collectors and local streets (US)

The standard toll fee is charged for every kilometre driven on every road subclass. Toll-eligible vehicles with onboard units installed get rebates for the kilometres driven on motorways and

¹¹ US terms: see Roess, Roger P. et.al.: Traffic Engineering, Third Edition, N.J, 2004

through roads. The level of rebate may differ for motorways and through roads. The instances of the road subclasses are depicted as tracks and are distinguished by the mean traffic which they carry.

Does the topology of the road system have a significant influence on the interactive behaviour of the toll-eligible vehicles? Experience shows that in segments where roads join or where the width of a road is reduced congestion is more frequent. Unfortunately no comprehensive data about these congestion sources exist for the relevant real world. Therefore it seems prudent to do without them in the model. The tracks of the road subclasses are implemented in the simulation view window as simple ribbons of equal width without interchanges and arranged parallel to the x-axis, coloured in different shades of grey. Each ribbon is divided into equal-sized patches.

The instances of the vehicle classes move along the ribbons from left to right. Because the simulation view window uses the wrap-around feature, vehicles leaving at the right edge reappear on the left.

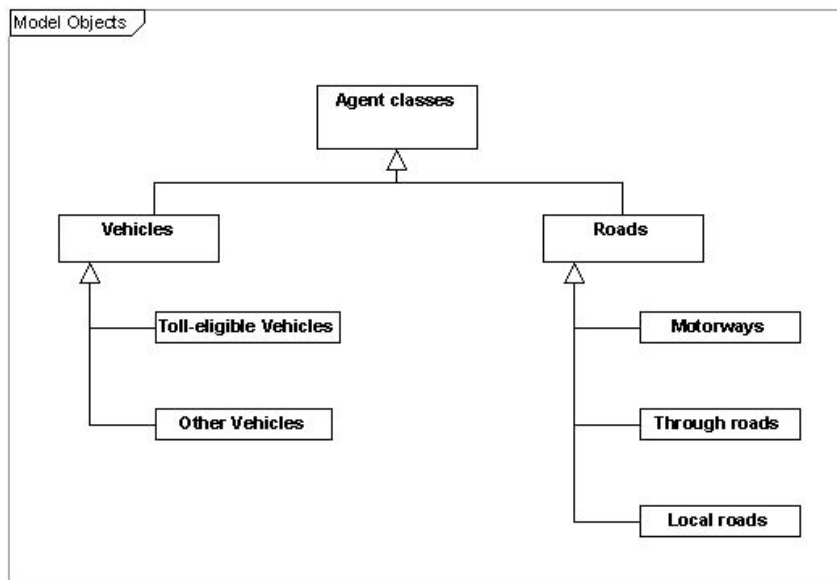


Fig. 2: Agent classes

Step 6: Scales

In the preceding section the agent classes and subclasses and their instances were defined at least in principle. In a next step the cardinality of the instances must be decided upon. For this the granularity of the simulation view window is crucial. Tests show that a 12" screen with a resolution of 1024 x 786 patches consisting of 9.5 pixels allows one to see sufficient detail. With patches of this size a simulation view window of 101 columns and 41 rows is reasonable and leaves enough space on the screen for other necessary input and output fields.

Each of the 41 rows of the simulation view window is an instance of a subclass of road. In this model the local roads are only feeders for the higher-ranking roads. It is assumed that congestion on them will not influence the transfer of traffic between road subclasses. All toll-eligible vehicles have their home and goal locations in patches which are an instance of local roads. As will be explained later, 500 or 600 instances of toll-eligible vehicles will be used. To visualise these in the set-up phase and later during the simulation, 11 rows are sufficient for the instances of the local roads.

Every instance of road “owns” 101 columns or patches. It proved useful to interpret a patch as a stretch of road which a vehicle driving at 50 km/h passes in one minute. The length of a patch is equivalent to 833 m in the real world. The length of an instance is about 85 km, but due to the wrap-around effect it is equivalent to about 100 km in the real world.

To prevent mapping errors in congestion behaviour the same scale must be applied for the cardinality of the motorway and the instances of the through-road subclass. As stated earlier, in that part of the real world which is relevant to this study there are 10,000 km of through roads with 20,000 lane kilometres and 2,000 km of motorways with 4,000 lane kilometres ¹². There are 30 rows available in the simulation view window to depict these. That means 25 rows can be assigned to the through road instances and 5 to the motorway instances. The scale for the kilometres is approximately 1 : 10.

Finally one must determine the number of instances of toll-eligible vehicles. As said earlier, 500 vehicle icons can be recognised sufficiently well in the simulation view window. With 500 instances to be represented, the scale of toll-eligible vehicle instances is 1 : 80.

Real-world observations can best be applied in a modelled state when they are not distorted by scale. This would mean that the scale should be 1 : 1. In our model this situation occurs when a toll-eligible vehicle gets caught in congestion or drives in a queue. Its behaviour depends on the distance to the vehicle ahead when this vehicle starts breaking or accelerating. This distance is about 50 m. The scale for this distance in the model, denoted X, in combination with the other scales must be 1 : 1, as formulated as a rule at the beginning of this chapter. Thus:

$$X \times (1 : 10) / (1 : 80) \Rightarrow (1 : 1).$$

The scale 8 : 1 fulfils the condition. This means that in the model a vehicle starts to react when its distance to the next vehicle ahead is 50×8 or about half the length of a patch, which is 416 m.

¹² It is assumed for purpose of the model that the toll-eligible vehicles drive only on one lane of the motorways.

Step 7: Movements and states

NetLogo, like most other multi-agent simulation packages, is a discreet simulation software, not an event-driven one. Every simulation run consists of equally long time steps. In this model one step is equivalent to one minute in reality.

At the start of a simulation the subclasses of the roads are created in the simulation view window. Next the icons for the instances of the other vehicles are put on the tracks of the through-road and motorway subclasses. The density distribution of the vehicles over the tracks must be similar to that observed in reality, as explained in greater detail later. The tracks of the through roads are placed in ascending order of vehicle density in the simulation view window, the tracks of the motorways in descending order. This makes visual analysis of congestion on the through roads and motorways easier. The distribution of the instances of the other vehicles in the tracks is uniform random. Finally the icons for the instances of the toll-eligible vehicles are put on the tracks of the local roads in a uniform random fashion. The icons are still blue because the decision as to which vehicle will have an onboard unit installed will be made later. Fig. 3 shows the simulation view window after set-up.




Fig. 3: Simulation view window immediately after set-up

At every simulation time-step every vehicle performs one of the moves described as follows:

First every vehicle checks whether its current speed has already reached the maximum speed allowed in the current track for that vehicle type; if not, it increases its speed. After this the vehicle moves along its track to the right for as many parts of a patch as it can drive in one minute, depending on its current speed. If it overshoots the length of route segment planned for the current road class, the vehicle immediately moves vertically up or down to the next track, as defined in its route plan. There it checks its speed and eventually adapts it.

If the vehicle is in a track on a through road or motorway, it also checks whether any other vehicles are within its reaction distance. If this is the case, it slows down its speed to that of the slowest vehicle and thereby starts a queue or adds itself to an already existing queue. A vehicle in a queue moves at the same speed as its predecessor. A queue starts disbanding when the leading vehicle has covered the length of route segment planned for the current road class and has moved to a track in another road class.

A toll-eligible vehicle which has come to the end of its current journey receives new values for its route-related attributes and starts its next journey at the next simulation step.

Fig. 4 shows the states a toll-eligible vehicle can be in and the transitions between states.

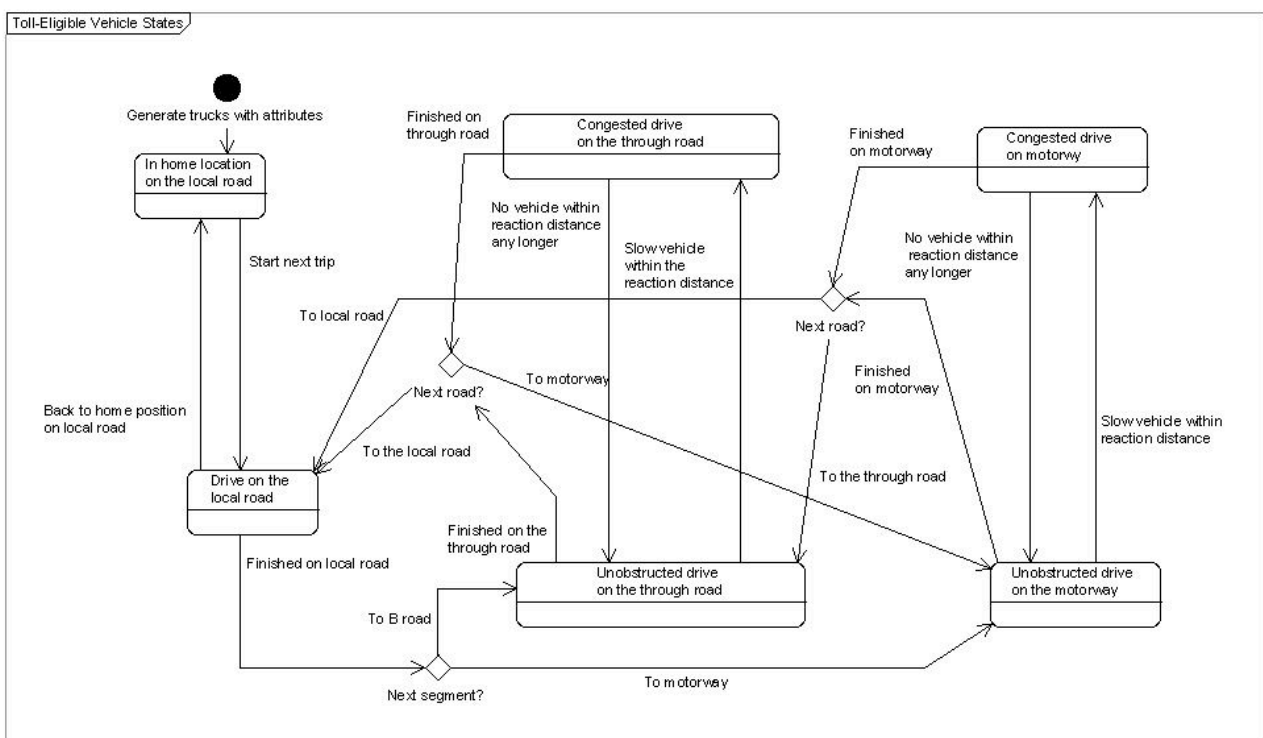


Fig. 4: The states of a toll-eligible vehicle

The role of the other vehicles is to help model realistic congestion situations. Therefore no specific journeys need be generated for them. They move along in their tracks and only carry out the following manoeuvres if they encounter a queue.

At the beginning of every simulation step the vehicle checks if any other vehicles are within its reaction distance in its current track or in the immediately neighbouring tracks to the left or right. If not and if not yet driving at its permitted maximum speed, the vehicle increases speed and moves on. If yes, it lowers its speed to that of the slowest vehicle ahead and moves on. If there is no other vehicle in one of the immediately neighbouring tracks within its reaction distance, it jumps to that track, adapts its speed appropriately and moves on. This means vehicles can overtake other vehicles.

At the start of a simulation step vehicles which are not in their home track jump to that track unless hindered by other vehicles there.

Fig. 5 shows the states which an instance of the class of other vehicles can be in and the transitions between states.

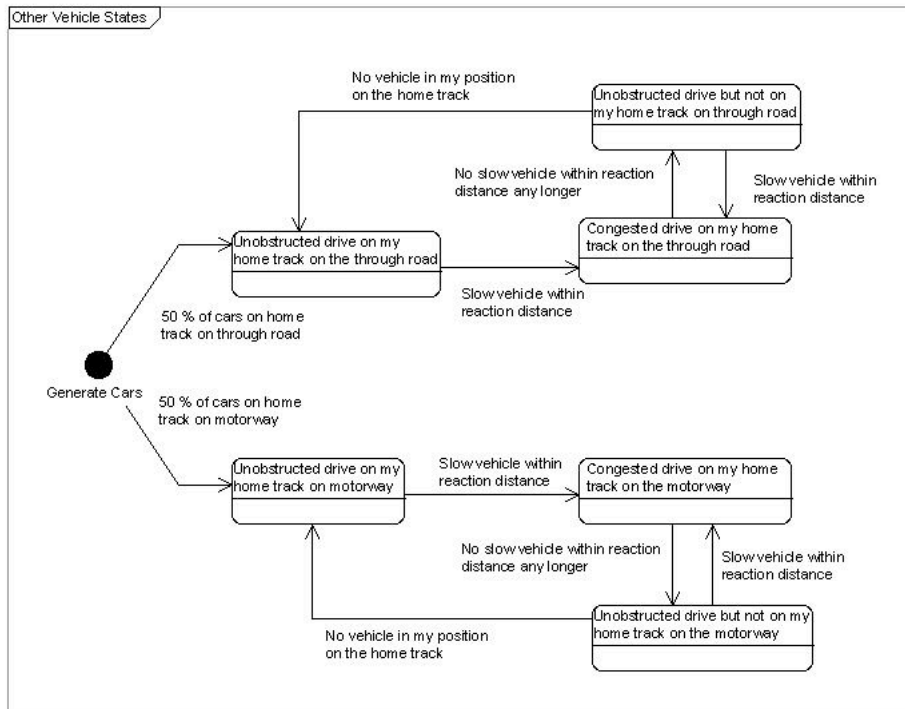


Fig. 5: The states of other vehicles (cars).

Fig. 6 shows an example snapshot of the simulation view window after about 1000 simulation steps with some queues on the motorway tracks and also some on the through road tracks.




Fig. 6: Example of the simulation view window after about 1000 simulation steps

Step 8: Behaviour functions

In this simulation model three groups of behaviour functions are applied:

- at the set-up of a simulation run: functions to generate the instances of the agent classes and the values of their static attributes,
- at the start of a new journey of a toll-eligible vehicle: a function which determines the values of the attributes characterising the route of the next journey and the behaviour during the journey when specific states are encountered,
- during the simulation run: functions which govern the interaction between instances depending on their current states.

At the set-up of a simulation run the values of global attributes like the standard toll fee or the price of an onboard-unit including installation costs are determined. Then the tracks of the road subclasses (co-ordinates and colours) and the maximum speed allowed on them are generated. In addition distribution functions are activated which assign to every track on the through roads and motorways the percentage of vehicles starting there. This reverse J-shaped function is derived from the relevant data of the real world. For example it defines that 22 % of the vehicles will start on the first track of the motorway and only 7 % on the fifth. Fig. 7 shows the distribution for through roads.




Fig. 7: Distribution of vehicles over the tracks of the through roads

After this the instances of other vehicles and their attributes are generated. The number of instances depends on the scenario being investigated. Half of the instances are randomly distributed over the through roads and the other half over the motorways. The vehicle distribution over the tracks of the road classes is the same as described above. These tracks are the vehicles' home tracks.

Next the instances of the toll-eligible vehicles with their attributes are generated and in a uniformly random distribution placed on the patches of the tracks of the local roads. The number of instances depends on the scenario being investigated. A very important attribute of the vehicles is their mean journey length. The journey length distribution is derived from the relevant data of the real world and has a monomodal shape which is skewed to the right with a mean value of 60 km. Another important attribute indicates whether the vehicle has an onboard unit installed or not. The decision whether to install an onboard unit or not takes into account the values of rebates, the price of the onboard unit and other realistic factors concerning the journey pattern of the toll-eligible vehicle, as is explained below. The decision is postponed to the end of the start-up phase of the simulation run.

At the start of every journey of a toll-eligible vehicle the total length of the standard journey and the lengths of its route segments on the three road subclasses must be generated. The total journey length is calculated by means of a stochastically distributed J-shaped function with a mean equal to the value of the vehicle's mean journey length as generated earlier. Using uniformly random numbers and with the maximum values for the segment lengths set as global attributes, the segment lengths on the local and through roads of the standard route are then generated. The segment length remaining for the motorways is simply the difference between the total length of the journey and the sum of the segments on the local and the through roads.

For vehicles which have an onboard unit installed a check is performed to find out if the standard route can be optimised by taking advantage of rebates. This must take into account the possible savings from driving on an optimised route as well as the delays due to congestion which may be encountered on this route. As many kilometres as possible are transferred to higher-ranking roads to optimise the route. Minor detours which make the optimised route slightly longer than the standard route are accepted. Intelligent guesses are used to set the minimum segment length for the road subclasses.

If many other vehicles optimise their routes, the probability of congestion on the higher-ranking roads will increase. This is related to problems which are the domain of the Minority Game¹³. Contrary to the approach of the Minority Game, there is a predefined solution tactic in this case. Modelling the behaviour procedure which maps the dispatcher's estimation of the risk of delays due to

¹³ A NetLogo version of the Minority Game is implemented by Daniel Stouffer; see Wilensky, U. (2004). NetLogo Minority Game. <http://ccl.northwestern.edu/netlogo/models/MinorityGame>. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL

congestion depends largely on episodic observations and the transfer of experiences¹⁴ from other traffic situations. The decisive quantity is the standardised relative time advantage between driving on the motorway vs. on the through road¹⁵. Its expected value for the next journey is calculated by weighting and combining the experience which the vehicle gathered from its last journey, the exponentially smoothed experiences from all other journeys of the vehicle, information about the current situation on the motorways and through roads, and the exponentially smoothed information from past periods about the motorways and through roads. The weights are one of the means to calibrate and fine-tune the model, as will be discussed later.

Looking at experiences with the truck toll system in Germany, one can assume that dispatchers base their decision to deviate from the standard route mainly on the possible savings in toll charges. Only if these do not seem worthwhile, do they consider the consequences of delays and change the standard route by transferring kilometres from the motorways to the through roads or even onto the local roads. This behaviour is modelled by defining a threshold value for the savings, which is then applied to the rebates and which must be surpassed before the dispatcher deems it cost-effective to optimise the journey route.

The route which will be stored in the vehicle attributes is defined by five values: km on local roads “up”, km on through roads “up”, km on motorways, km on through roads “down” and km on local roads “down”. If the total journey length is very small, then segments on the through roads or the motorways may become zero.

Step 9: Time regime of the simulation

Progress during the simulation and the states at the end of a run may be disturbed by singular start-up behaviour and later by strong oscillations. These effects can mask the real results or produce distorted findings. A properly designed time regime for the simulation runs will eliminate or at least attenuate them.

This model shows a distinct start-up regime due to the simultaneous generation of all toll-eligible vehicles in their home locations. Trial runs demonstrate that the start-up behaviour subsides when each vehicle has finished on average 10 journeys. At this point the decision will be made in which vehicle an onboard-unit is to be installed.

Vehicles with onboard units may optimise rebate payments by transferring traffic to the higher-ranking road classes. But then the probability of congestion and delays increases, which makes op-

¹⁴ See Underwood, Geoffrey(Ed): Traffic and Transport Psychology: Theory and Applications, Amsterdam, 2005

¹⁵ standardised relative time advantage between driving on the motorway vs. on the through road =
(speed on motorway / speed on through road) / (effective top speed on motorway / effective top speed on through road)

timising less attractive. The oscillation caused by this feed-back loop never ceases but its amplitude becomes relatively stable after a certain number of simulation steps. Only then should the sampling of simulation results start. In this case-study sampling starts after each vehicle has finished 30 journeys on average. To limit computer power run time, the simulation is stopped after each toll-eligible vehicle has finished 50 journeys on average.

Simulation models apply random number generators. To eliminate the effect of their seed values, 10 runs are performed in this model for each scenario and the results registered. These are condensed using standard statistical methods and their spread visualised.

Step 10: Debugging and validation

The model was implemented with NetLogo. Although debugging is supported by NetLogo features, for instance by the possibility to single out one agent, put a spotlight on it and follow its movement in the simulation view window during a simulation run, it is still a demanding exercise. Software tools for multi-agent simulation on personal computers do not handle the activities of the agent in parallel but transform them into sequences, sometimes in an idiosyncratic way. Debugging needs a thorough understanding of what the code really does. The best way to gain this insight is to take some time “playing” with the program until one really knows how the agents behave and interact and which piece of code triggers what. All the while one should aim to improve the code by simplifying or amending it until it is parsimonious and perfectly transparent.

The bug-free functioning of a multi-agent simulation program does not automatically mean that it maps the relevant part of the real world correctly. It must be validated. The standard validation procedure would be to compare the results from the simulation program with known data about the real world. This does not really work with multi-agent models because at best only global data are available for the real world and not at all, or not in sufficient detail, for a sufficient number of agents. What is known about the behaviour of real agents is often only episodic, biased and not representative. As a surrogate for the missing real data a plausible scenario about the behaviour of agents and the whole system was created. This scenario is based on information from professionals working with freight companies, toll operators and the authority in charge of monitoring freight traffic on roads. Missing aspects were supplemented with “intelligent” predictions and guesses.

The program was walked through, executed and adapted until it was recognisably no longer false or implausible; i.e., the agents behaved as expected in all their possible states and the relationships of the result variables showed correct signs and magnitudes.

Step 11: Calibrating

Before a validated multi-agent simulation program can be applied to produce quantitative predictions for alternative scenarios it must be calibrated. Here calibrating means finding through a trial-and-error process values for parameters which are not directly observable and measurable in the real world. In this case-study two sets of parameters are used. One set consists of the weights in the behaviour procedure which estimates the standardised relative time advantage between driving on the motorway as opposed to on the through road. The other is the threshold which allows the dispatcher to decide if the route for the next journey of a toll-eligible vehicle should be optimised or not.

To calibrate the weights they must be varied until a calibration benchmark is satisfied. Here the proportion of drivers of toll-eligible vehicles who, before the launch of the toll system, perceived serious delays on motorways due to congestion outside main holiday weeks was used as a calibration benchmark. The benchmark value is 5 %. To find the combination of values for the calibration parameters which best satisfies the benchmark, their behaviour space¹⁶ with the weights as dimensions was generated and evaluated.

When calibrating a multi-agent simulation model sometimes no meaningful combination of values for the parameters can be found which will satisfy the benchmark. Normally the problem can be solved with an adaptation of the behaviour procedure of which the parameters are a part. Sometimes more far-reaching changes are required.

In this case-study the calibration process uncovered a minor error in the program. After its correction and re-validation of the model, values for the parameters were found which give a reasonably close approximation to the benchmark value even when the parameter values were varied slightly.

Step 12: Alternative scenarios and behaviour space

A series of alternative scenarios was executed with the validated and calibrated simulation program. As input data those explained above under 4) were used:

- Rebate schemata:
 - Schema 1: through roads 0 %, motorways 0 %
 - Schema 2: through roads 5 %, motorways 5 %
 - Schema 3: through roads 5 %, motorways 10 %
- Number of toll-eligible vehicles:
 - 500

¹⁶ NetLogo has a feature called "BehaviorSpace". It allows the alternatives and the number of runs to be defined and executes them in batch mode. After each run the input data and the result data are collected in a matrix.

600

- Number of other vehicles:
100
200

Each of the 18 combinations of input data was run 10 times.

Findings of the case study

Results

About 10 hours of computer time were needed to execute the runs for the alternatives. The resulting matrix with the input and result data for the alternatives was further processed. Fig. 8 shows a selection of the result variables in condensed form.

Rebate schema	Toll eligible vehicles	Other vehicles	Mean % km added to motorways	Mean % km added to through roads	Mean % vehicles on motorway
1	500	100	-3.0	3.4	31.3
		200	-13.9	13.9	30.5
	600	100	-7.0	7.8	31.5
		200	-20.7	19.6	30.2
2	500	100	1.1	2.6	33.9
		200	-6.2	10.8	33.5
	600	100	-2.1	6.4	34.3
		200	-16.1	19.8	31.9
3	500	100	17.2	-8.8	43.6
		200	7.1	5.2	41.1
	600	100	10.0	0.5	43.0
		200	0.2	12.0	39.1

Fig. 8: Results of the simulation

The values are the arithmetical means of the data collected by the BehaviourSpace feature of NetLogo during the simulation steps between the point when on average every toll-eligible vehicle finishes 30 journeys to the point when on average every vehicle finishes 50 journeys. Because 10 runs

were evaluated for every result variable, the arithmetical mean of these 10 values then had to be calculated.

The effect of the random numbers on the mean percentage of toll vehicles driving on motorways is shown as “box and whisker plots”¹⁷ (see Fig. 9). The boxes in the middle indicate quartiles and medians. The lines show the largest/smallest observations that fall within a distance of 1.5 times the box sizes. The x-axis is composed of the alternatives. The first 4 columns stand for the alternative where no rebates on the standard toll fee are considered, the next 4 for the alternative with 5 % rebate for all kilometres driven on through roads and motorways, the last 4 for the alternative with 5 % rebate for all kilometres driven on through roads and 10 % for all kilometres driven on motorways.

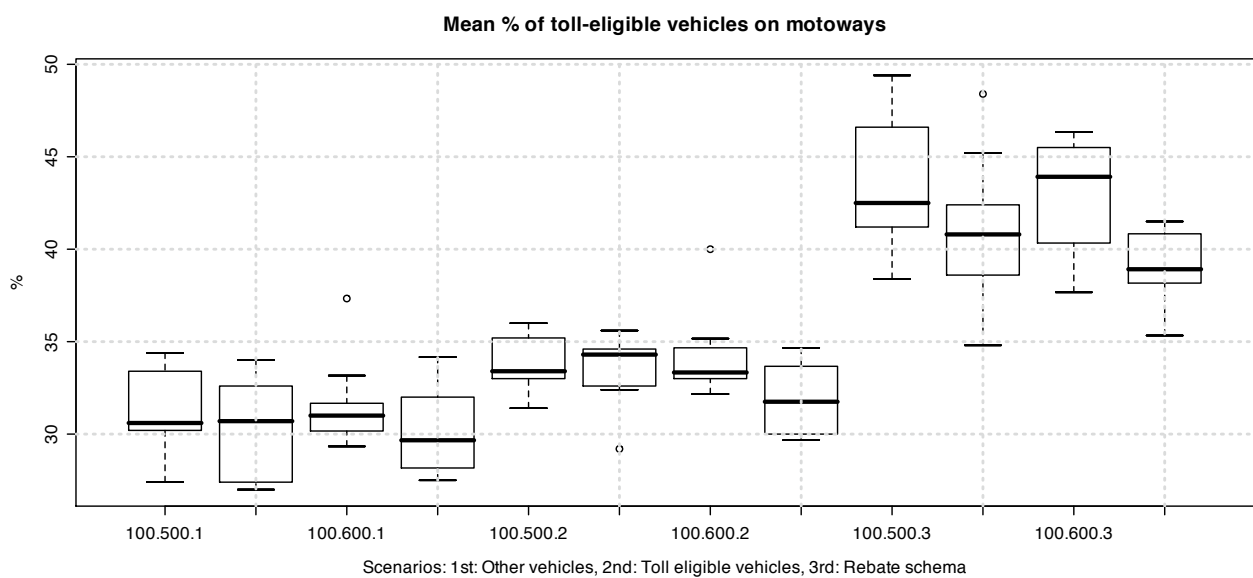


Fig. 9: Box plot based on 10 simulation runs

Fig. 10 shows in one column the minimum and maximum mean values of the rebates effectively gained by the vehicles with installed onboard units for the 3 rebate schemata. These rebates are slightly smaller than the possible maximum, which would be 5 % for the second rebate schema and about 7.5 % for the third rebate schema. The deviations are caused by the detour kilometres and the transfer of traffic to lower-ranking roads to avoid congestion.

The other column shows the percentage of onboard units installed in the toll-eligible vehicles. As long as rebates are small, as in the second rebate schema, onboard units make sense from an economic point of view, and under the assumptions made here, only for those vehicles which on average drive long journeys.

¹⁷ Dalgaard, Peter: Introductory Statistics with R, New York, 2002, p. 65

Rebate schema	Mean effective rebates %	Onboard units installed %
1	0	0
2	4.2 - 4.4	28 - 31
3	6.7 - 6.8	94 - 95

Fig. 10: Rebates and onboard units.

Conclusions

The values of the result variables look reasonable in their tendency and magnitude. The higher the rebates, the more traffic is transferred to the higher-ranking roads. Increasing the number of vehicles halts this trend but only to a small extent.

Before quantitative conclusions can be drawn from the results the influences of the simplifications which have been made in the simulation model must be discussed. In the model a toll-eligible vehicle starts its next journey immediately after the previous one has ended. In reality the start of the next journey is always delayed and follows a pattern over the hours of the day, days of the week and over the weeks of a year. As a consequence of these patterns the traffic load and the probability of congestion change over time. When the peak traffic hours are over, congestion usually dissolves quickly if not caused by accidents. These effects are not included in the model. Therefore the congestion-related results are to some extent distorted, as discussed later.

Including the start pattern of the vehicles in the model would have increased the complexity of the code and the number of periods per simulation run. As it is, the model uses about 1500 steps in a run to collect values for the result variables. To map the daily traffic pattern 1440 steps would be needed for one day (in the model one simulation step equals one minute in real time). To achieve statistically stable results a simulation run would have to cover at least 10 days. Hence, about 10 times more simulation steps than used in the study now would be needed. This is beyond the capability of the available computer.

The percentage of toll-eligible vehicles driving on motorways increases more than linearly when the rebates are increased. While in the second rebate schema the proportion of vehicles on motorways increases only 0.6 percentage points for 1 percentage point of rebate increase compared to the scenario without rebates, for the third rebate schema the increase is 1.8 percentage points. In reality this number will be slightly higher because the model overestimates the delays caused by congestion. On the other hand congestion caused by accidents or bad weather conditions are not included.

The proportion of vehicles driving on motorways is not very sensitive to a change in the number of vehicles. A 20 % increase in toll-eligible vehicles combined with a 100 % increase in other vehicles reduces the proportion of vehicles driving on motorways by 3 percentage points for the second rebate schema and by 3.7 percentage points for the third. Hence, an increase in other vehicles, for instance during holidays, has a greater effect than an increase in toll-eligible vehicles. Due to the simplifications in the model the reduction in the proportion of vehicles driving on motorways is overestimated.

The results of the multi-agent simulation model allow the road owners to draw helpful conclusions when considering the introduction of a comprehensive truck toll system. Nevertheless, road owners should bear in mind that the model results are slightly negatively biased due to the simplifications discussed above. Negative here means that the results look less advantageous for the road owners than they would be in reality. A few of the conclusions derived from the results presented above are spelled out as follows:

- Only when the rebates offered are higher than 5 % and strongly differentiated between through roads and motorways will there be a considerable traffic shift. Rebates below 5 % have too small an impact on the traffic structure, do not reduce the externality costs significantly and therefore are not worth the effort. To find the rebate schema with the best cost/benefit ratio the externality costs and other relevant influences must be assessed and weighted, which is beyond this study.
- Those bottlenecks on the motorways which are especially sensitive to the number of other vehicles should be eliminated before the launch of the toll system. Otherwise delays caused by congestion and experienced by the dispatchers may unduly hinder the shift of toll-eligible vehicles to the motorways.
- The shift of traffic to the motorways is not very sensitive to changes in the number of toll-eligible vehicles. An increase in the number of toll-eligible vehicles over the next years will not substantially increase pressure on the road owners to expand the capacity of their motorways.

A global corporate model was developed for the comprehensive toll system. There the traffic shifts between through roads and motorways are modelled by behaviour equations. To find appropriate mathematical functions for these behaviour equations and their parameters, econometric methods were used which evaluated the data generated by runs of the multi-agent simulation model.

It is hoped that this study has demonstrated that multi-agent simulation models are a useful device for dealing with real-world problems and has shown how to proceed in such circumstances. Some-

times, as in this case, it would seem that multi-agent simulation is the best approach available. But the study should also have demonstrated that developing such models is a demanding task, requiring detailed understanding of the relevant reality and meticulous work. Models based on global or aggregated variables definitely take less time to implement. Maybe this is the reason why in recent years many agent-based demonstration models have been announced and offered for downloading on the internet but real-world based ones have rarely become widely known.