

## MAINTAINING A DISTRIBUTED SYMBIOTIC RELATIONSHIP USING DELEGATE MULTIAGENT SYSTEMS

Rutger Claes  
Tom Holvoet

DistriNet Labs  
Katholieke Universiteit Leuven  
Celestijnenlaan 200A, 3000 Leuven, BELGIUM

### ABSTRACT

Online simulation of traffic can assist route guidance systems by predicting problems such as congestion. Accurate predictions require accurate status information about vehicles - the fact that the vehicles are distributed over large-scale road infrastructure makes this particularly challenging. Embedding the online simulation in the road infrastructure - by distributing it across road side computing infrastructure - is a partial solution, but also adds additional complexity to the symbiotic relationship between online simulation and the physical system. In this paper we describe an approach that uses delegate MultiAgent Systems to reduce the complexity of such symbiotic relationships. Experimental results in a prototype implementation of the route guidance mechanisms show that the approach is feasible and leads to a proactive route guidance mechanisms with the potential of outperforming current state of the practice non-proactive routing mechanisms.

### 1 INTRODUCTION

The past decade the use of onboard route guidance systems has risen drastically. Some of these onboard devices are capable of receiving traffic updates and, if desired, divert their users to avoid delays such as traffic jams, accidents or other obstacles. These devices can be considered reactive. They only take into account existing problems, trying to divert drivers to avoid a further increase in congestion and minimize the travel time based on known delays.

Another approach is to avoid the congestion altogether. This is a proactive approach in which drivers are guided to avoid forming traffic jams, thus further reducing drivers travel time. Proactive traffic guidance relies on forecast information. Using this forecast information traffic jams can be predicted and anticipated, allowing route guidance mechanisms to divert drivers proactively. In their paper (Kaufman, Smith, and Wunderlich 1991) Kaufman, Smith and Wunderlich describe a system capable of providing forecast traffic information to drivers and show using experimental results that a proactive route guidance mechanisms can decrease travel times.

One of the mechanisms that can generate the forecast information required by such a proactive approach is online, or symbiotic, simulation as described in (Ayd et al. 2008). Symbiotic simulations are simulations operating in close relationship with the physical system they represent. The term *symbiotic* refers to symbiosis in biology where it is used to describe close and often long-term interactions between different biological species. A symbiotic simulation benefits from the physical system it simulates by receiving continuous real-world information that can be used to make the simulation more accurate. The physical system, on the other hand, also benefits from the simulation because it can use the outcome of so-called “what-if” simulations to guide decision making. This relationship is demonstrated in Figure 1.

Traffic, by nature, is a distributed system: vehicles are scattered across the road infrastructure. When traffic is the subject of an online simulation, the physical system shown in Figures 1 is distributed. Consequently, maintaining the symbiotic relationship between physical system and simulation becomes more difficult. It is no longer a 1 : 1 relationship, it becomes an  $n$  : 1 relationship as depicted in Figure 2b.

This also has consequences for the opposite side of the relationship. The online simulation has to process information from the physical system. If the physical system is distributed and the online simulation remains centralized, a bottleneck will be formed. Keeping the online simulation centralized can have consequences regarding the scalability, reliability and robustness of the approach.

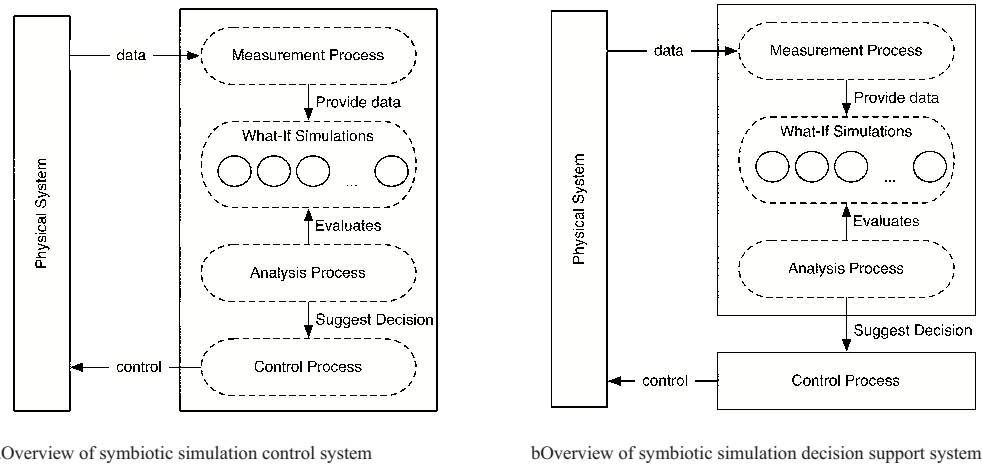


Figure 1: High level overview of two of the closed loop symbiotic simulation systems proposed in (Aydt et al. 2008).




Figure 2: Relationships with different cardinalities. The fully distributed  $n : m$  relationship in Figure 2c is challenging to maintain.

This paper describes a fully distributed proactive route guidance system. Distributing the online simulation and maintaining the symbiotic relationship between both the distributed physical system and the distributed online simulation is challenging. Instead of a  $n : 1$  relationship, the result is a  $n : m$  relationship as shown in Figure 2c. The symbiotic relationship itself becomes distributed. Our approach combines two techniques to tackle these problems and achieve a fully distributed proactive route guidance mechanism. The first technique is the use of embedded online simulations. The second technique is the use of delegate MultiAgent Systems to maintain the symbiotic relationship by aggregating and propagating information.

When distributing the online simulation, the approach presented here draws inspiration from the approach outlined in (Hunter et al. 2009): distributing the simulation based on domain entities such as vehicles and roads. Our approach distributes the online simulation across road side computer infrastructure. In several countries such road side infrastructure is planned for as part of road pricing schemes. The challenge of synchronizing these distributed simulations and providing input information for these simulations are tackled throughout this paper.

To reduce the complexity of maintaining the distributed symbiotic relationship we apply some of the patterns of delegate MultiAgent Systems as described in (Holvoet, Weyns, and Valckenaers 2009). Some of these patterns, such as the use of delegate MultiAgent Systems have been applied successfully to the traffic domain before (Weyns, Holvoet, and Helleboogh 2007).

While the results shown in this paper are preliminary, and larger and more diverse experiments still need to be conducted, we believe the results obtained with our prototypical application are promising.

The remainder of this paper is structured as follows. The paper starts with a discussion on related work involving proactive route guidance and embedded online simulations. The next section outlines the approach proposed in this paper. The paper continues with an evaluation of the proposed approach. Finally, we end with a conclusion in which

we reflect back on the use of delegate MultiAgent Systems in maintaining the symbiotic relationship between physical system and distributed embedded online simulations.

## 2 RELATED WORK

There exists a body of work describing proactive route guidance using forecast information. In this section we describe the two existing approaches that closely resemble our approach and inspired it, namely the work on *Link Travel Time Predictions* presented in (Wunderlich, Kaufman, and Smith 2000) and the work on *Ad Hoc Distributed Simulations* described in (Fujimoto et al. 2007).

### 2.1 Link Travel Time Predictions

The route guidance system, called *Link Travel Time Predictions* (Wunderlich, Kaufman, and Smith 2000) uses forecast instead of real-time data to help vehicles make routing decisions. This system is able to provide vehicles with forecast data based on a central traffic simulation. In their paper, Wunderlich et al. offer the following definition of anticipatory route guidance:

Anticipatory route guidance can be defined as paths calculated from a set of predicted future link travel times, which when disseminated to drivers cause the same set of predicted link travel times to be realized by vehicles in the network.

Like the approach by Wunderlich et al. the approach described by Kaufman, Smith, and Wunderlich (1991) is similar to the approach presented in this paper, vehicles are given forecast traffic information and use this information to anticipate future traffic when making routing decisions.

The experiments conducted by Kaufman et al. reveal a decrease in average travel time for the anticipatory vehicles. The other vehicles also benefit from the presence of anticipatory vehicles in the network. A lot of attention is given to the penetration rate of the anticipatory vehicles. It is shown that as the penetration rate increases the benefit the anticipatory vehicles have compared to the other vehicles diminishes. Partly because non-anticipatory vehicles benefit from more anticipatory traffic, but also because the average travel time of the anticipatory traffic starts to increase.

In the approach described by Wunderlich et al., all vehicles send their travel intentions to one central Independent Service Provider or ISP. At this central ISP, a dynamic traffic assignment process is used to predict future traffic densities. This ISP then distributes these link travel time predictions to all vehicles with a subscription, thus providing them with forecast information that would lead to anticipatory routing decisions. Little attention is given to the specifics on how the intentions of the vehicles are gathered and especially how the forecast information is distributed back to the vehicles.

The authors in (Wunderlich, Kaufman, and Smith 2000) make a distinction between centralized and decentralized route guidance systems. The differentiating factor between a centralized and decentralized system is where the route is calculated. If this is done centrally, and then gets dispatched to the participating vehicles, the system is said to be centralized. If the vehicles make their own planning, albeit based on information originating in a centralized simulation, the system is said to be decentralized.

This distinction between centralized and decentralized route guidance systems can be expressed in terms of the definition of a *Symbiotic Simulation Decision Support System* defined in (Aydt et al. 2008). In such a *Symbiotic Simulation Decision Support System*, depicted in Figure 1b, a physical system provides information to an online simulation. This online simulation can provide decision support to a *Decision Maker*, which influences the physical system. The distinction between centralized and decentralized systems as defined in (Wunderlich, Kaufman, and Smith 2000) can be reduced to a decision based on the location of the *Decision Maker*. If every vehicle has its own *Decision Maker*, choosing routes based on information stemming from the online simulation the mechanism is distributed. If the *Decision Maker* is centralized and proposed routes for all vehicles - more resembling the *Symbiotic Simulation Control System* shown in Figure 1a described by Aydt et al. - the mechanism is labeled as centralized.

For the approach we present in this paper we only consider a system to be decentralized when there are no centralized components. The use of a centralized simulation required in the approach described by Wunderlich, Kaufman, and Smith (2000) introduces a centralized component and many problems typical for centralized solutions. Distributing the route guidance system would offer the potential of increased robustness, responsiveness and reliability.

### 2.2 Embedded Ad Hoc Distributed Simulations

In (Fujimoto, Hunter, Sirichoke, Palekar, Kim, and Suh 2007), the authors present a novel way of distributing an online simulation. *Ad Hoc Distributed Simulations* are simulations residing onboard the vehicles they simulate. Every participating vehicle simulates its own area of interest. The vehicles synchronize these simulations using aggregates and can rollback if the discrepancies become too large. Ad Hoc Distributed Simulations are defined in the paper as:

An ad hoc distributed simulation is a collection of autonomous on-line simulations brought together to model an operational system.

The work by Fujimoto, Hunter, et al. focusses on distributing the online simulations. There is little attention given to closing the loop, that is, using the information generated by the online simulations to control the vehicles in the physical system. In this respect, this approach can be catalogued as an Open-loop Symbiotic Simulation System according to (Aydt et al. 2008). The forecast information generated by ad hoc distributed simulations can be used for proactive vehicle routing or traffic monitoring as stated in the paper by Fujimoto et al.

The concept of *embedded simulations* is further emphasized in (Hunter et al. 2009). In this paper the authors envision

embedding on-line distributed simulations operating within transportation network elements (participating vehicles, roadside cabinets, etc.) in an attempt to create a transportation management and prediction system distributed across in-vehicle computer systems, roadside computers, and traffic management centers. In this approach participating vehicles play an active role in the monitoring and prediction of future states of the transportation network.

The focus in both of these papers remains on distributing the online simulation across the participating vehicles. These simulations include information obtained from road side sensor networks and traffic management centers. As vehicles are able to simulate multiple possible future trajectories, this approach is able to do the “what if” simulations referred to in the symbiotic simulation models shown in Figure 1.

### 3 DISTRIBUTING BOTH SIDES OF THE SYMBIOTIC RELATIONSHIP AND MAINTAINING IT

In this section we outline our approach. Traffic is an open and distributed environment, vehicles are distributed across the road infrastructure and continuously enter and exit the system. In our approach, we model vehicles and roads as *situated agents*. A situated agent is an agent embedded in a virtual environment. Situated agents interact with this virtual environment. Together, these agents form a situated MultiAgent system. The desired functionality of the system, providing the vehicles with forecast information so they can make better routing decisions, comes not from these individual agents, but rather from their cooperation and interactions. Both the *situated agents* and the *virtual environment* patterns are described by Weyns in (Weyns 2009).

The situated MultiAgent approach we present here has two distinct types of agents. RoadAgent’s to represent the road elements and VehicleAgent’s to represent vehicles participating in the system. RoadAgent’s are deployed on road side computing infrastructure. This requirement on the presence of a road side computing infrastructure, together with vehicle to infrastructure communication is realistic because of the dependence of many other initiatives such as road pricing schemes, on this infrastructure. VehicleAgent’s are deployed onboard vehicles.

VehicleAgent’s know the destination of their vehicle. They ask the virtual environment to offer them routing options to reach that destination. In return, the VehicleAgent’s inform the RoadAgent’s of the route they have chosen. This route choice represents the intention of the VehicleAgent. Based on this information on the vehicles intended routes, the RoadAgent’s estimate future traffic intensities on the road element they represent and offer better route options when asked by the VehicleAgent’s.

The proactive vehicle routing mechanism described here is based on the work presented in (Holvoet and Valckenaers 2007, Weyns, Holvoet, and Helleboogh 2007), where the use of delegate MultiAgent Systems as coordination mechanism in *coordination and control problems* and anticipatory vehicle routing is introduced.

#### 3.1 Distributing the online simulation across road side infrastructure using embedded simulations

In our approach, RoadAgent’s host an online simulation and use this simulation to estimate future link travel times and offer route choices to the VehicleAgent’s. The model used in this simulation to predict travel times could be learned by the embedded simulation, this is facilitated by the fact that the simulations are hosted by RoadAgent’s and thus are located close to the actual road they represent. The future traffic demand is provided by the VehicleAgent’s, just as in the reservation-based case.

In the approach presented by Weyns et al., the RoadAgent’s treated the information received from the VehicleAgent’s as reservations. The RoadAgent discretizes the road element it represents and keep track of all reservations. It then uses these reservations to estimate future traffic intensities and link travel times. This information is then used to advice VehicleAgent’s on their route.

Using the reservations to provide possible routes for the VehicleAgent’s stems from other application domains such as manufacturing control in which the anticipatory vehicle routing approach by Weyns et al. was first applied. In most of these domains using reservations proves to be a good solution. In traffic, however, it is not.

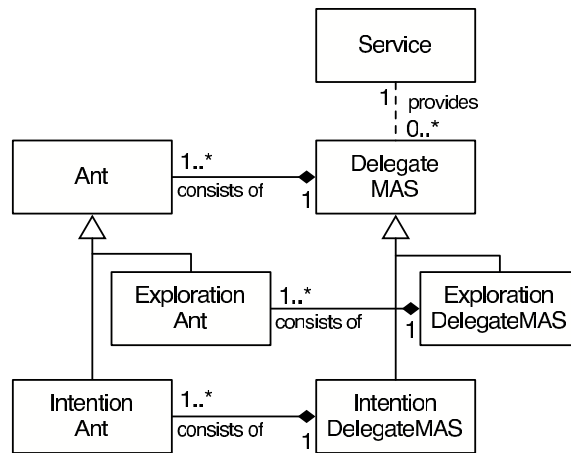


Figure 3: Schematic representation of delegate MAS. Together, virtual ants form delegate MultiAgent Systems and can provide a service to other agents.

It is hard to maintain reservations on a discretized road element when confronted with vehicles of different size, driving at different speeds. The use of reservations is very problematic when not all vehicles in the system participate in the coordination scheme as mentioned by Dresner and Stone in (Dresner and Stone 2007). Non-participating vehicles would “steal” reservation slots from the participating vehicles, thus making all future reservations invalid. Even if all vehicles would participate, because traffic is a non-controlled environment, enforcing reservations becomes impossible. Because of these problems, this paper abandons the use of reservations in the traffic domain. The use of embedded simulations as defined by Fujimoto and Hunter proves to be much more promising.

The approach we take here differs from that presented by Fujimoto. There, the authors also embed the online simulation in the environment. The difference, however, is that in the approach presented in this paper, we place the simulations on the road infrastructure, while in their paper, the simulations are distributed over the vehicles. As such, the ad hoc distributed simulations described by Fujimoto et al. can be seen as a longitudinal simulation of the vehicles trajectory. In contrast, placing the simulations on the road side infrastructure offers a cross sectional view on the trajectories of all vehicles planning to traverse the road element.

Using roadside embedded simulations thus distributes the online simulation. Because both sides of the symbiotic relationship are distributed across the environment - the vehicles are, by nature, scattered across the road, the simulations are embedded and are placed next to road elements - maintaining the relationship is a challenge.

The simulation model used in the embedded simulations are relatively simple. Using a model describing the average speed of a vehicle dependent on the number of vehicles that are entering the road or are already on the road we can use a small event based simulation to simulate a small piece of road. The model that describes the vehicles average speed can be fine tuned at runtime.

### 3.2 Using Delegate MAS to maintain the symbiotic relationship

In (Weyns, Holvoet, and Helleboogh 2007), the communication between VehicleAgent’s and RoadAgent’s is modeled using virtual ants in a delegate MultiAgent System. The use of such biologically inspired coordination mechanisms has been described in (Dorigo, Maniezzo, and Colomi 1996) and has been described as a coordination pattern in (Holvoet, Weyns, and Valckenaers 2009) (Figure 3).

The propagation of VehicleAgent’s intentions and the aggregation of forecast information produced by the online simulations is delegated to a separate MultiAgent System, hence called a delegate MultiAgent System. Such a delegate MultiAgent system consists of smaller, mobile agents. These smaller agents often use pheromones to communicate with other agents and are consequently referred to as “ants”. In our proposed approach, both directions of the symbiotic relationship depicted in Figure 1, the physical system providing data to the online simulation and the simulation providing decision support to the system, are handled by delegate MultiAgent Systems consisting of ants as shown in Figure 4.

The use of ants and pheromones in traffic routing has been described in (Narzt et al. 2007), (Ando et al. 2006) and (Tatomir and Rothkrantz 2006). In this paper the use of delegate MultiAgent Systems, and consequently the use of ants, is considered to be a vital part of the solution.




Figure 4: Intention and exploration ants maintaining both directions of the symbiotic relationship between vehicles and distributed embedded online simulation.

### 3.2.1 Exploration of the environment based on traffic forecast information

In the approach described in this paper *VehicleAgent*'s explore the environment using a delegate MultiAgent System. This is in line with reservation based approach described in previous works in literature (Weyns, Holvoet, and Helleboogh 2007). The ants operating in this delegate MAS are called "exploration ants". Exploration ants explore the virtual environment. The virtual environment defines a graph based on the actual road infrastructure it represents.

The exploration ants search this graph. They start from the position of their origin *VehicleAgent* and search for paths to the destination. While searching for a path, they continuously pose "what if" questions on behalf of their *VehicleAgent* to all *RoadAgent*'s they encounter. These "what if" questions trigger "what if" simulations, similar to those of Figure 1. Such a "what if" simulation is scoped to the road element the *RoadAgent* represents and the time at which the vehicle will pass this road element.

The result of such a "what if" question is memorized by the exploration ant and used to specify an arrival time when posing the next "what if" question. On reaching its destination, the exploration ant reports the path it traversed back to its origin *VehicleAgent*. The information the ant reports back consists of the possible route leading to the vehicles destination annotated with forecast link travel times indicating how long it would take the vehicle to reach its destination when choosing this route.

This mechanism of using ants to explore the virtual environment is a suitable abstraction even when using online simulations instead of reservations. Because the ants remember the answer to the previous "what if" question, the outcome of an online simulation is propagated along the ants path to initialize the next online simulation. Hence, the exploration ants act as synchronization between the distributed online simulations residing on the individual road elements. A second function of the exploration ants is the aggregation of all simulation outcomes along a possible path on behalf of the *VehicleAgent*. The exploration ants are responsible for maintaining one direction of the symbiotic link between simulation and vehicles.

### 3.2.2 Propagating vehicles intentions to the online simulations

Once a vehicle chooses the path it intends to follow, based on the information gathered by the exploration ants, it has to propagate this information back into virtual environment. This information from the vehicles is propagated through the virtual environment and made available to the *RoadAgents* using another type of delegate MultiAgent System.

The *VehicleAgent* delegates this functionality to a delegate MultiAgent System consisting of "intention ants". These ants traverse the path the vehicle intends to follow. Whenever they encounter a *RoadAgent*, the intention ants informs the *RoadAgent* of the vehicles pending visit by depositing an amount of "virtual pheromone". By aggregating and interpreting the deposited amounts of pheromones, *RoadAgent*'s know the amount of vehicles intending on driving on their road segment in the near future and can use this information to initiate the "what if" simulations used to provide information to visiting exploration ants.

We refer to the information deposited by intention ants as pheromones because they evaporate over time and because the information is deposited at a particular location. The evaporation property of the information guarantees freshness of the data. If the information is not reinforced - or refreshed - at regular intervals, it evaporates and eventually will be discarded. When vehicles change their intentions, choose another path to reach their destination, the information on their previous intention evaporate over time. Without the need for the *VehicleAgent* or the ants to update the old intention. To maintain the information about their travel plans, however, *VehicleAgent*'s need to continuously dispatch intention ants.

Intention ants, like exploration ants, also report back to the *VehicleAgent* that dispatched them. The information they report back is identical to that reported back by the exploration ants. This information can be used to reconsider the intention if a faster route is found or if the current intention becomes blocked due to an incident.

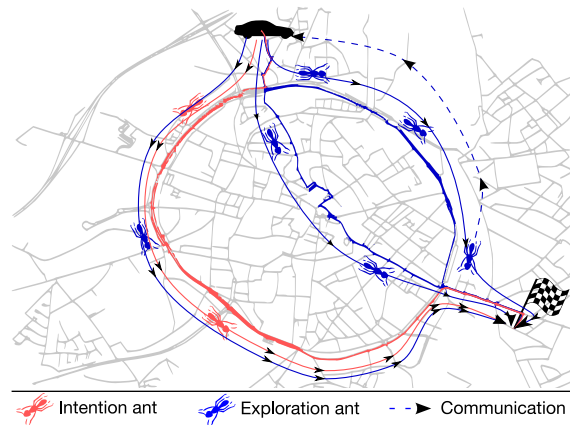


Figure 5: Ants exploring the environment of the city of Leuven. Exploration ants explore the environment by aggregating forecast information generated by embedded online simulations distributed across the road infrastructure. Based on this information, the `VehicleAgent` chooses a route it intends to follow and informs the `RoadAgent`'s along this path using intention ants.

In (Weyns, Holvoet, and Helleboogh 2007) intention ants are used to make reservations. As described earlier, the approach presented in this paper does not use the notion of reservations, having replaced it with online distributed simulations. The use of intention ants to propagate `VehicleAgent`'s intentions to the `RoadAgent`'s is still applicable when using online simulation instead of reservations.

#### 4 PROTOTYPICAL ILLUSTRATION

We have implemented a proactive route guidance system based on the techniques identified in Section 3. Based on simulations we evaluate the applicability and performance of our implemented coordination mechanism by comparing it with a non-proactive, reactive route guidance system.

##### 4.1 Scenario description

All simulations use the city of Leuven and its surroundings shown in Figure 6 as traffic infrastructure. We model this traffic infrastructure based on OpenStreetMap data<sup>1</sup>. The information offered by OpenStreetMap allows to model the roads realistically as it includes information on road characteristics and junction characteristics.

During the simulation vehicles are spawned on the traffic infrastructure. Every vehicle has a predefined destination it must drive to. The rate at which the vehicles spawn depends on the type of scenario we want to simulate. A first scenario type inserts vehicles at a steady input rate of  $n$  vehicles per second. A second scenario type inserts vehicles at a varying input rate. In this scenario periods of low traffic intensity and high traffic intensity follow each other. This type of scenario can also be described by the average vehicle input rate  $n$ . We refer to this last scenario type as the dynamic input rate scenario.

The origin and destination of the vehicles is determined at random. The probability distribution used in the random selection of locations favors trips starting outside the beltway of Leuven and ending on the other side of the city, also outside the beltway. Using this probability distribution we generate a synthetic population of 28800 vehicles. This population is then injected at different rates into the simulation for each of the evaluated coordination mechanisms. This approach ensures all test results originates from the same origin destination combinations, injected in the simulation in the same order.

To evaluate the proactive route guidance system, we compare it with a non-proactive route guidance mechanism based on current state of practice technologies. This alternative route guidance system is modeled after onboard satellite navigation systems enhanced with traffic information broadcasts. This routing strategy is also reactive, it does not anticipate traffic intensities. It is, however, no longer static. The onboard satellite navigation system receives information about existing traffic problems and is able to suggest alternatives to the driver. This routing strategy is referenced as *TMC* in the remainder of this paper. TMC is the acronym for Traffic Message Channel, the technology that delivers the traffic updates to the onboard satellite navigation systems over the radio.

<sup>1</sup><http://www.openstreetmap.org>

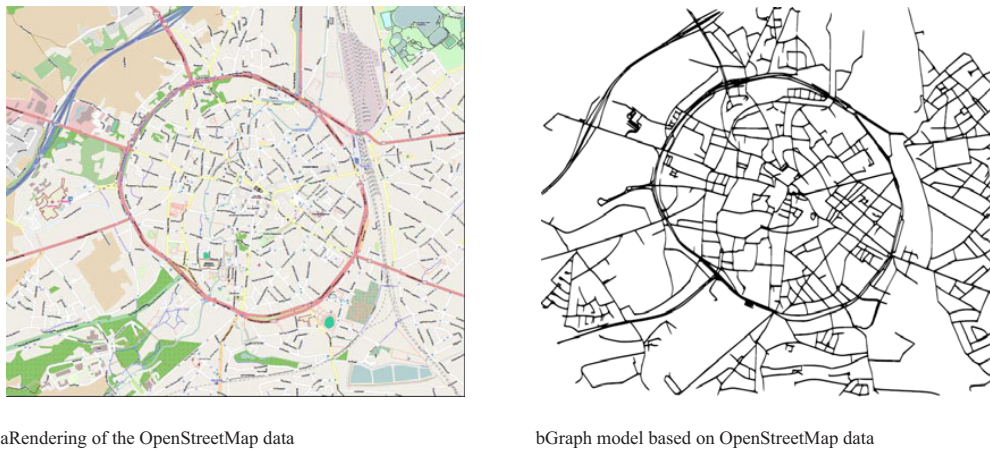


Figure 6: Map of simulated region. Figure (a) on the left shows a rendering of the OpenStreetMap data. Figure (b) shows the graph with over 1600 roads used in the micro-simulation.

These experiments are simulations of distributed simulations and their interaction with other systems and are aimed at minimizing travel times, not maximizing performance. As such, details on the hardware and the software used to run the simulation are of little relevance. The distributed simulation is not simulated on a hardware level. For the simulation of the distributed simulation performance is not an issue as it is for evaluation purposes only.

## 4.2 Results

In the scenario described above, our delegate MultiAgent Systems based coordination mechanism managed to guide all vehicles to their destination. The travel times for all vehicles were measured and allow us to compare the efficiency of our distributed approach based on anticipatory data with the traditional TMC aided routing approach. By calculating the mean travel time for vehicles using the TMC aided approach ( $mean_{tmc}$ ) and for those using the ant inspired approach ( $mean_{ant}$ ) we are able to compare the mean travel times of both approaches in Figure 7.

In Figure 7 the gain, defined as  $100\% \times (1 - mean_{ant}/mean_{tmc})$  is shown for various input loads and for both dynamic and static input rates. The graph shows that, as the input rate increases, a proactive routing strategy outperforms a non-proactive routing strategy.

The  $mean_{tmc}$  and  $mean_{ant}$  numbers only give a limited indication on the overall performance of both routing strategies. However, they do show that on average our distributed approach is on par with the traditional centralized approach.

## 5 CONCLUSION

In this paper we have described a proactive traffic route guidance mechanism using online embedded simulation distributed across the road infrastructure and a delegate MultiAgent Systems based symbiotic relationship. Based on the implementation of this mechanisms and the experiments performed with the implementation we can conclude two things, namely (1) Proactive traffic guidance mechanisms have the potential to outperform reactive traffic guidance mechanisms, even when using distributed embedded online simulations; and (2) Maintaining the symbiotic relationship between a distributed physical system and a distributed online simulation using patterns of delegate MultiAgent systems is possible.

The fact that proactive traffic guidance can outperform purely reactive traffic guidance systems is not surprising. When a coordination mechanism has access to traffic forecast data, it should be able to get some advantage out of this data. Obtaining the forecast data based on the intentions of all participating vehicles and distributing this information back to the individual vehicles are challenging when developing a distributed proactive traffic guidance system.

A central online simulation component can be avoided by distributing the online simulations over the traffic infrastructure, similar to the approach described in (Hunter et al. 2009). The monitoring and learning from the physical system by the online simulation is facilitated by the fact that the online simulation is deployed in close proximity to the actual road element it simulates.

Maintaining the relationship between physical system and online simulation when both sides are distributed over the same environment can be challenging. Using some of the patterns of delegate MultiAgent Systems helps in reducing



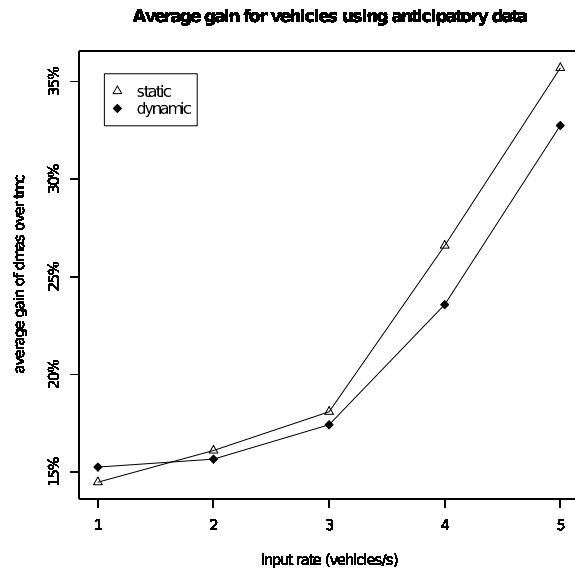


Figure 7: Gain in average trip duration obtained by using forecast instead of real-time data.

the complexity in both the online simulation nodes, part of the RoadAgent's, and the nodes of the physical system, represented by the VehicleAgent's. The use of pheromones and stigmergy enables the VehicleAgent's to change their intentions without having to revisit the abandoned paths, thus preserving the freshness of the information on which the online simulations operates.

#### ACKNOWLEDGMENTS

This research was funded by the IWT - SBO project 'MASE' (project no. 060823) and by the Interuniversity Attraction Poles Programme Belgian State, Belgian Science Policy, and by the Research Fund K.U.Leuven.

#### REFERENCES

- Ando, Y., O. Masutani, H. Sasaki, H. Iwasaki, Y. Fukazawa, and S. Honiden. 2006. Pheromone model: Application to traffic congestion prediction. *Engineering Self-Organising Systems*:182–196.
- Aydt, H., S. Turner, W. Cai, and M. Low. 2008. Symbiotic simulation systems: An extended definition motivated by symbiosis in biology. In *Proceedings of the 22nd Workshop on Principles of Advanced and Distributed Simulation*, 109–116. IEEE Computer Society.
- Dorigo, M., V. Maniezzo, and A. Colomi. 1996. Ant system: optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 26 (1): 29–41.
- Dresner, K., and P. Stone. 2007. Sharing the road: autonomous vehicles meet human drivers. In *IJCAI'07: Proceedings of the 20th international joint conference on Artificial intelligence*, 1263–1268. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Fujimoto, R., M. Hunter, J. Sirichoke, M. Palekar, H. Kim, and W. Suh. 2007. Ad hoc distributed simulations. In *PADS '07: Proceedings of the 21st International Workshop on Principles of Advanced and Distributed Simulation*, 15–24. Washington, DC, USA: IEEE Computer Society.
- Holvoet, T., and P. Valckenaers. 2007. Exploiting the environment for coordinating agent intentions. *Environments for Multi-Agent Systems III, LNAI 4389*:51–66.
- Holvoet, T., D. Weyns, and P. Valckenaers. 2009. Patterns of delegate mas. In *2009 Third IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, 1–9. IEEE.
- Hunter, M., J. Sirichoke, R. Fujimoto, and Y.-L. Huang. 2009. Embedded ad hoc distributed simulation for transportation system monitoring and control. In *Proceedings of the 2009 INFORMS Simulation Society Research Workshop*, ed. L. H. Lee, M. E. Kuhl, J. F. Fowler, and S. Robinson, Volume 1, 13–17.
- Kaufman, D., R. Smith, and K. Wunderlich. 1991. An iterative routing/assignment method for anticipatory real-time route guidance. In *Vehicle Navigation and Information Systems Conference, 1991*, Volume 2.

- Narzt, W., G. Pomberger, U. Wilflingseder, O. Seimel, D. Kolb, J. Wieghardt, H. Hortner, and R. Haring. 2007, 30 2007-Oct. 3. Self-organization in traffic networks by digital pheromones. *Intelligent Transportation Systems Conference, 2007. ITSC 2007. IEEE*:490–495.
- Tatomir, B., and L. Rothkrantz. 2006, Sept.. Hierarchical routing in traffic using swarm-intelligence. *Intelligent Transportation Systems Conference, 2006. ITSC '06. IEEE*:230–235.
- Weyns, D. 2009. A Pattern Language for Multi-Agent Systems. In *Proceedings of the 2009 Joint Working IEEE/IFIP Conference on Software Architecture European Conference on Software Architecture (WICSA/ECSA 2009)*.
- Weyns, D., T. Holvoet, and A. Helleboogh. Sept. 30 2007-Oct. 3 2007. Anticipatory vehicle routing using delegate multi-agent systems. *Intelligent Transportation Systems Conference, 2007. ITSC 2007. IEEE*:87–93.
- Wunderlich, K., D. Kaufman, and R. Smith. 2000. Link travel time prediction for decentralized route guidance architectures. *IEEE Transactions on Intelligent Transportation Systems* 1 (1): 4–14.

#### **AUTHOR BIOGRAPHIES**

**RUTGER CLAES** is a PhD student in DistriNet labs at the Department of Computer Science, Katholieke Universiteit Leuven. He graduated as a master in applied sciences and engineering at the Katholieke Universiteit Leuven. His current research interests include Multi-Agent systems, Multi-Agent based simulations and decentralized coordination mechanisms. His email address is [<rutger.claes@cs.kuleuven.be>](mailto:rutger.claes@cs.kuleuven.be).

**TOM HOLVOET** is a professor in the Department of Computer Science, Katholieke Universiteit Leuven. His research interests include software engineering of decentralized and multi-agent systems, coordination, software architecture, and autonomic computing. He received his PhD in computer science from the Katholieke Universiteit Leuven in 1997. Contact him at [<tom.holvoet@cs.kuleuven.be>](mailto:tom.holvoet@cs.kuleuven.be).