

Technical report 09-052

Multi-Agent Control of Urban Networks – Algorithm and Case Study*

R. T. van Katwijk and B. De Schutter

To cite this work, please refer to the published version:

R. T. van Katwijk and B. De Schutter, “Multi-agent control of urban networks – Algorithm and case study,” *Proceedings of the 16th ITS World Congress*, Stockholm, Sweden, Sept. 2009. Paper 3017.

Delft Center for Systems and Control
Delft University of Technology
Mekelweg 2, 2628 CD Delft
The Netherlands
phone: +31-15-278.24.73 (secretary)
URL: <https://www.dcsc.tudelft.nl>

* This report can also be downloaded via <https://dpub.eu/09-052>

MULTI-AGENT CONTROL OF URBAN NETWORKS ALGORITHM AND CASE STUDY

R.T. van Katwijk, Netherlands Organization for Applied Scientific Research, TNO, Delft,
The Netherlands, phone: +31 (0)15 2697473, email: ronald.vankatwijk@tno.nl
B. De Schutter, Delft Center for Systems and Control & Maritime and Transport Department,
Delft University of Technology

ABSTRACT

As a result of increasing congestion and the sharper constraints on the traffic system with respect to throughput, safety, and air quality, the number of traffic control instruments will increase both in number and in heterogeneity. As more and more traffic control instruments are installed to promote the flow of traffic the probability increases that conflicts will arise or that coordination opportunities are lost when traffic control instruments are applied in the same area. Traffic control instruments can thus no longer be considered separately, but have to be considered as part of a larger network. By modeling the separate instruments as intelligent agents, the actions of the individual instruments can be coordinated. This paper illustrates the benefits of multi-agent coordination and defines a procedure and a new movement-based look-ahead traffic-adaptive control algorithm through which coordination between traffic control instruments can be achieved.

INTRODUCTION

To improve traffic flow, safety, and air quality road authorities have installed many traffic control instruments, such as traffic lights, ramp metering installations, dynamic route information panels and “motorway signaling” (for autonomous speed control and lane closure commands).

The majority of the installed traffic control instruments functions fully autonomously and has been tuned to attain a predetermined local objective. Local disturbances in the traffic flow are addressed without having to call upon other traffic control instruments or a higher traffic management level and can thus be dealt with more timely and efficiently. As more and more instruments are deployed, chances are that instruments will be deployed in each other’s region of influence, resulting in interference whenever the actions of the individual instruments are not coordinated. Especially in urban areas, intersections are located so closely together that intersections directly influence the traffic demand for neighboring intersections. So in order to guarantee the effective functioning of the network as a whole it is required that the deployed traffic control instruments are coordinated.

In the beginning of the eighties online urban network traffic management became a reality as a result of the development of SCOOT and SCATS [1, 2]. In those early years the control parameters of the traffic lights could only be adapted with a limited frequency and for a limited amount of control parameters. Since then these network traffic management systems have evolved over a number of generations. All these systems employ a centralized or top-down approach to network traffic management.

A system that determines the control parameters from within a traffic management center is *theoretically* able to determine control parameters that are optimal. Although this approach

is very appealing it just is not always possible to do this efficiently and effectively, which is largely due to the amount of data involved and the computational complexity of the problem. By allowing the individual instruments to coordinate their actions based on the information they receive from sensors and from each other, traffic control instruments can be coordinated more often and more accurately than can be done by a traffic operator. This paper therefore proposes a multi-agent approach that is able to mitigate the shortcomings of a centralized approach and serve as a complement to the aforementioned centralized systems.

This paper is organized as follows. First, we describe our multi-agent approach to the network control problem. Next, the procedure we propose to coordinate the actions of the agents is described. Subsequently, the new movement-based look-ahead traffic-adaptive control algorithm we propose to control the intersections in the network is described. Next, the dynamics of the multi-agent approach are explained using an illustrative example. Finally, the results of a simulation study are presented which compared the multi-agent approach with traffic-actuated control.

MULTI-AGENT APPROACH

One way to more effectively handle disturbances is to make the traffic control instruments more intelligent and to have them deal with the intricacies of configuring the traffic control instruments to the situation at hand. The traffic management center can then focus on traffic management instead of on traffic control, as the latter is fully dealt with by the traffic control instruments.

A distributed approach requires that the control problem can be separated into multiple loosely coupled subproblems, such that the composition of the solutions to all the subproblems approaches the solution of the original control problem. Multi-agent systems can aid in the distribution of the problem and facilitate the coordination of the activities of the traffic control instruments when needed. The term agent is used to denote an intelligent actor that interacts with its environment by means of sensors and actuators. A multi-agent system in turn is a system comprised of several agents that together are capable to attain an objective that is difficult to attain by a single agent or monolithic system.

Agent technology can make an important contribution to traffic management since the abstract concepts and ideas as used in the domain of agents and multi agent systems easily translate to the traffic management domain [3]. Most multi-agent systems are explicitly designed so that cooperation between the individual agents is obvious. This expertise with the principles of coordination in a heterogeneous environment of agents in general is directly beneficial to the development of these principles for traffic management in particular.

In order to create a more intelligent traffic control instrument it is necessary to make the current traffic control instruments that thus far have been operating purely locally more susceptible to the interest of its neighbors. The agent controller we propose uses an improved look-ahead traffic-adaptive control algorithm that is capable to incorporate downstream and upstream traffic conditions in its decision making. This makes it possible for the agent controller to coordinate its actions with neighboring agent controllers. The procedure to coordinate the actions of neighboring traffic control agents is described in the next section. The improved look-ahead traffic-adaptive control algorithm is described in the subsequent section.

COORDINATION PROCEDURE

The idea of coordination between traffic control instruments is very old [4]. In urban areas, with a large number of signalized intersections, traffic engineers often face the task of coordinating multiple subsequent intersections in order to allow platoons of vehicles to move through several signalized intersections without stopping. The movement of a platoon of vehicles through several signalized intersections is referred to as progression. To achieve progression, a timing relationship is developed between successive intersections such that vehicles, traveling at a predetermined speed, can pass through the green indications at successive signals. This limits the flexibility of the system.

The system thus created can be compared to the coordinated equivalent of the fixed-time control of a single intersection.

The efficiency of a progression scheme is largely dependent on the platoon ratio in the traffic stream. Platoon ratio is the fraction of arterial traffic that travels from the first intersection through the last intersection in the system. Signal performance will also depend on the amount of traffic on minor movements in two ways: (1) The traffic volume on the cross streets affects the percentage of traffic turning into the main street, and therefore affects the platoon ratio; and (2) a low volume on the minor movements allows the controller to spend more green time on coordinated movements.

Given the dynamics in traffic volume and platoon ratios a progression scheme is needed that is able to flexibly adapt to changes in volumes and platoon ratios. To achieve this flexibility a traffic control agent has to be aware of the effect of the actions of agents downstream and upstream on its own performance and the effects of its own actions on the performance of the agents downstream and upstream. In order for the traffic control agent to be able to ascertain whether delay prevented upstream is annulled at the downstream intersection the traffic control agent also incorporates the expected downstream performance in its decision making. To this aim information is shared between intersections. This information is necessary in order to determine the delay a released vehicle will encounter as it approaches the downstream intersection. Whereas the actions of the downstream traffic control agent are considered as given, the own signal plan is the result of the optimization. As at first no information from neighboring control agents is available the signal plan is iteratively updated to incorporate the changes at neighboring control agents as they become available.

At each time step each agent determines its current state. It does so on the basis of the information it gathers from local detectors and from information received from upstream and downstream agents. At each time step each agent tries to optimize its performance by regulating its inflow and outflow. Requests are made to downstream agents by upstream agents concerning their intended outflow. When such a request is received the downstream agent determines the impact of the intended inflow on its performance. This impact, expressed in the form of a cost is communicated to the upstream agent that made the request. The upstream agents can then decide whether the downstream costs outweigh their own costs, and make a decision regarding the outflow they want to realize. The coordination procedure is repeated until the agent no longer wishes to update its control settings.

The above procedure is summarized in the following high-level procedure:

In order to be able to incorporate downstream and upstream traffic conditions in the decision making of a control agent and to be able to do so iteratively in real-time we developed a new movement-based look-ahead traffic-adaptive control algorithm which is described in the following section.

Algorithm 1 High-level coordination procedure

```
1: loop
2:   while NOT (equilibrium or cycling) do
3:     update current state
4:     determine optimal control settings incorporating (a) downstream cost and (b) expected
       upstream inflow
5:     for all downstream agents do
6:       send intended outflow
7:       receive downstream cost of planned outflow
8:     end for
9:     for all upstream agents do
10:      receive planned inflow
11:      determine cost of upstream inflow given local and downstream cost
12:      send cost of upstream inflow
13:    end for
14:  end while
15: end loop
```

CONTROL ALGORITHM

There exist several control approaches for traffic signal control such as fixed-time, vehicle-actuated, or look-ahead traffic-adaptive control. The state-of-the-art in traffic signal control is currently formed by traffic-adaptive systems. However, in practice, the majority of the controlled intersections in most countries are still controlled by vehicle-actuated controllers (e.g., in The Netherlands about 85 %). Traffic-actuated controllers extend the green time for the current movement as long as there is demand for that movement. The total amount of green time that can be given to a movement is constrained by a maximum. A limitation of this type of controller is that its decision to extend the green for a movement does not take into account traffic present on the other, non-green, approaches and does not take into account traffic approaching the signal from further upstream.

Traffic-adaptive control differs from vehicle-actuated control because it can evaluate a set of feasible control decisions and make a decision that is optimal with respect to its control objective. A *look-ahead* traffic-adaptive controller additionally is capable of determining the optimal control decision on the basis of a longer-term analysis, and is often able to incorporate information from further upstream. This allows the look-ahead traffic-adaptive controller to make better decisions in the long run. “Regular” traffic-adaptive control can be considered to be short-sighted compared to look-ahead traffic-adaptive control.

Look-ahead traffic-adaptive control algorithms use a moving-horizon approach in which at each control step an optimal control sequence is determined for a given prediction horizon. Control decisions at each control step then involve whether or not to end or to initiate green for movements. This essentially results in an optimization problem involving a decision tree.

As the decision space has a discrete structure, the search for the optimal sequence of decisions corresponds to building a decision tree.

To illustrate the approach consider the decision tree depicted in Figure 1. The tree is rooted by the current state, labeled “now”. Four branches lead away from this state, each corresponding to the result of a control decision that can be taken at this state. The control decision made results in a new state from which again four control decisions can be made. Each decision takes

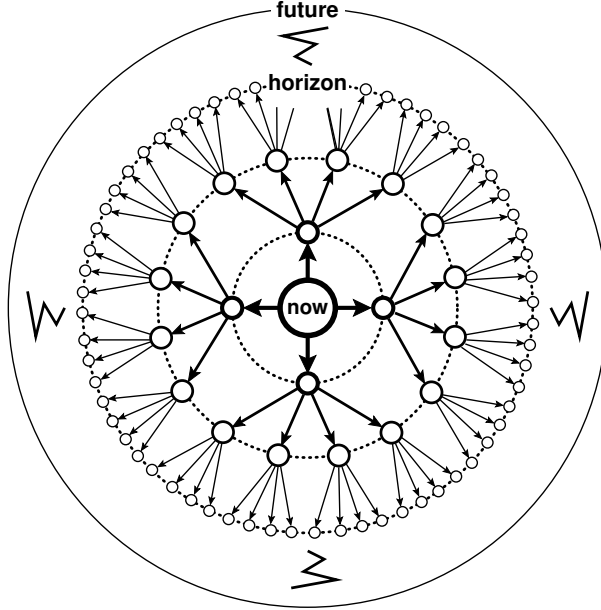


Figure 1: Decision tree.

us further into the future and deeper into the tree formed by the states resulting from each possible decision. Efficiency in searching the decision space is considered by the degree to which the entire tree will not have to be explored to find an optimal path.

Consider the intersection depicted in Figure 2(a). A *movement* corresponds to a stream of vehicles that could get green or red¹, such as, e.g., movement 11 in Figure 2(a) which represents the vehicles on the upper arm of the intersection that can drive straight ahead or turn left. Given the set of movements for an intersection, a *stage* is then a (fixed) assignment of red or green indications to each of the movements over a period of time.

For an intersection consisting of 12 separately controlled movements there are 111 possibilities to give green to those movements safely.

The state-of-the-art in look-ahead traffic-adaptive control still uses a stage-based approach since it is not possible to consider all possible combinations of movements as search space size grows exponentially with the number of possibilities considered. In a stage-based approach movements are assigned to so called stages. Movements within a stage get green (or red) for the entire duration of the stage. The look-ahead traffic-adaptive algorithm we propose allows a movement to switch, e.g., from green to red within a stage, at which time another non-conflicting movement can get green. Our approach, by contrast, does consider individual movements. As such it is more flexible than a stage-based approach as it allows green for signals in different stages to start sooner if the demand for all conflicting movements in the current stage has cleared. In order to gain the same level of flexibility with a stage-based approach as with a movement-based approach the search space has to be increased from 171^N to 111^N for an intersection with 12 separately controlled movements. In practice the number of separately controlled movements is often larger as a consequence of separately controlled pedestrian, bicycle and public transport movements.

In [7] several look-ahead traffic-adaptive control algorithms, such as PRODYN [8], OPAC [9], UTOPIA-SPOT [10], RHODES [11], and ALLONS-D [12], are assessed based on compu-

¹For the sake of simplicity of the exposition, we do not explicitly consider the yellow phase or the all-red phase, but for the time being we work with green and red only (see also [5, 6] for a more detailed explanation).

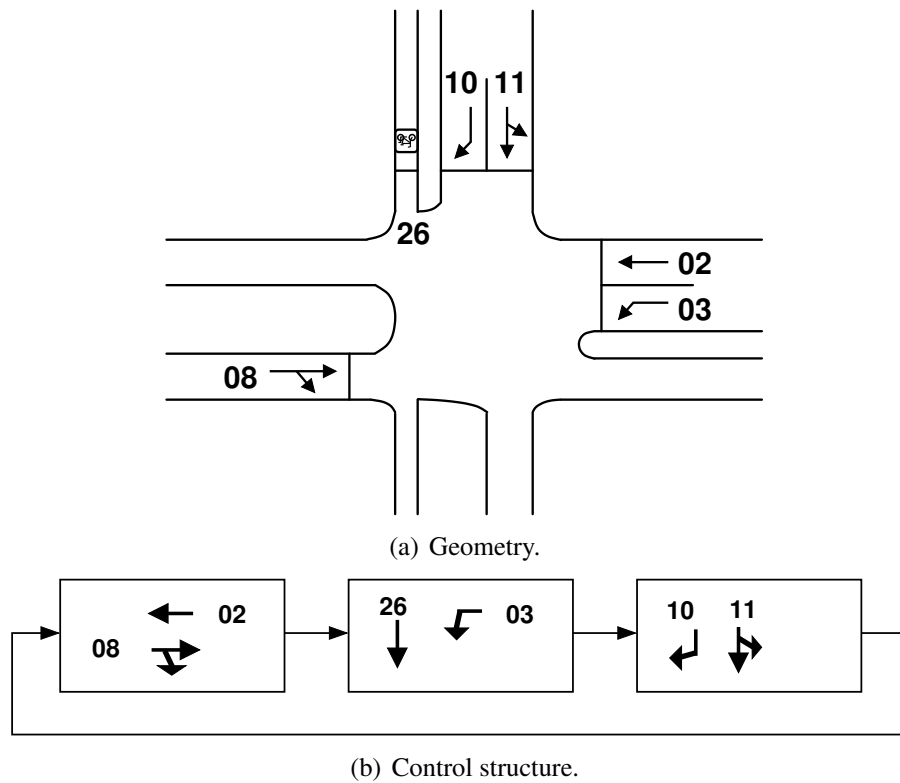


Figure 2: Intersection used as an example.

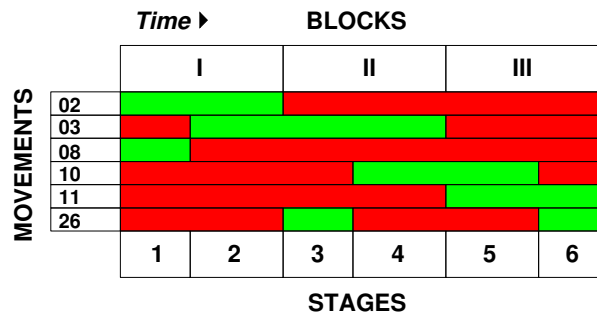


Figure 3: Number of blocks versus the number of stages for the intersection of Figure 2(a).

tational speed and on the quality of the results (in terms of vehicle delay). All these algorithms use a stage-based approach to traffic control as opposed to the more flexible movement-based approach employed by the state-of-the-art in vehicle-actuated control.

By adopting this block/movement-based approach (or movement-based approach for short) we can increase the number of green combinations without increasing the size of the search space.

This is illustrated in Figure 3. This figure shows a possible timing of green and red intervals for the movements given the structure of Figure 2(b). The period of time in which a movement gets green or red is denoted by a bar that is colored accordingly. The interval of time in which a block is active is depicted at the top side of the picture. The figure shows that as soon as green for movement 08 of block I has terminated, movement 03 of block II is allowed to get green. This is allowed since movement 03 has no conflict with the movements of block I that still get green (in this case, this is movement 02). Similarly, movement 10 of block III is allowed to advance to green as soon as the green phase for movement 26 has ended.

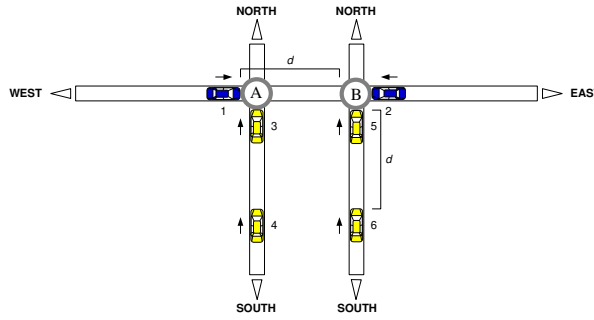


Figure 4: Geometry of the corridor used in the microscopic coordination example

In Figure 3 the advantage is illustrated of employing a movement-based approach over a stage-based approach for the intersection of Figure 2(a). As the movement-based approach allows movements to switch to green as soon as all conflicting movements have cleared, the number of green combinations possible with just three blocks would have required six distinct (non-maximal) stages if instead an equivalent stage-based approach had been applied. These six stages are depicted on the bottom of the picture. Optimizing over the horizon using blocks instead of stages significantly reduces the branching factor of the tree and thus significantly reduces the search space without making sacrifices with regard to the quality of the solution.

The search-algorithm employed is based on a dynamic programming formulation similar to the one found in [13]. A full description of the algorithm can be found in the first author's PHD thesis [5].

In the next section we will describe how a movement-based approach can be incorporated in a look-ahead traffic-adaptive control algorithm.

ILLUSTRATIVE EXAMPLE

The benefits and the dynamics of the multi-agent approach can be understood by considering the following example.

Figure 4 shows a corridor consisting of two intersections. The traffic demand consist of six vehicles, one eastbound approaching the western intersection (vehicle 1), one westbound approaching the eastern intersection (vehicle 2) and four northbound approaching the western (vehicles 3 and 4) and eastern (vehicles 5 and 6) intersections respectively. Blue vehicles are used to indicate the vehicles that travel from the west to the east or vice versa whereas the yellow vehicles indicate the vehicles that travel from the south to the north. Potential conflicts are therefore always between vehicles with a different color and never between vehicles of the same color. Note that vehicles 1 and 6 and vehicles 2 and 4 are at equal distance (d) from the western and the eastern intersection respectively. When two vehicles are predicted to arrive at an intersection and request opposing signal phases, a conflict is said to occur. Two of these conflicts are depicted in the time-space diagram of Figure 5(a). This diagram shows the distance to the downstream intersection(s) of vehicles 1, 3, and 6 as it evolves over time. The dashed lines mark the location of the intersections, whereas the other lines mark the location of vehicles 1, 3, and 6 as time progresses. The diagram shows that vehicle 1 has potential conflicts with vehicle 3 at the first downstream intersection and with vehicle 6 at the next downstream intersection. In the time-space diagram these events are marked with a star shape. In total there are four potential conflicts as Figure 5(a) only depicts the conflicts of vehicle 1 with vehicles 3 and 6. The potential conflicts that are not depicted are those of vehicle 2 with vehicles 5 and 4.

The intersection can deal with these potential conflicts in one of two ways, one suboptimal

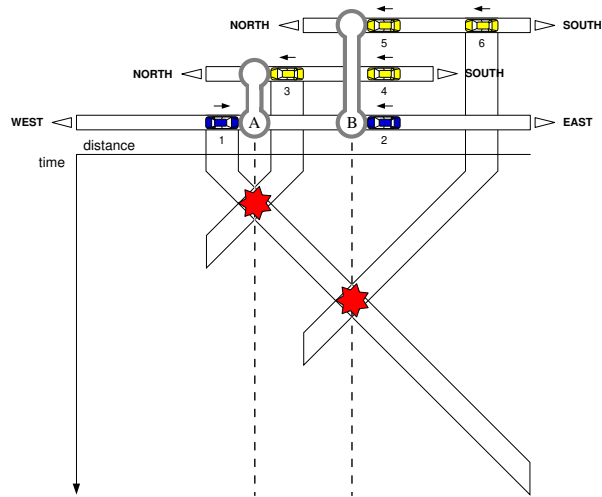
(as shown in Figure 5(b)) and one optimal (as shown in Figure 5(c)). In the suboptimal case preference is given to vehicle 1 over vehicle 3, whereas in the optimal case preference is given to vehicle 3 over vehicle 1. The vehicles that are not given preference to are stopped. These events are marked in the time-space diagrams by a stop symbol. From a local perspective both options can be considered equal as from a local perspective not all potential conflicts in the network are visible. As the western intersection is unaware of traffic approaching the eastern intersection it is unaware of the impact of its decision to first release vehicle 1. If the intersection releases this vehicle first, the vehicle will come in conflict with vehicle 6 on the eastern intersection. The coordination procedure can prevent suboptimal behavior at the local intersection. The process that results from applying the coordination procedure is described next.

Initially both controllers are unaware of what the situation is at nearby intersections. As the process of the western intersection mirrors that of the eastern intersection in the remainder only the former will be described in detail. The intersection starts out to observe the traffic state on its approaches. The intersection controller determines that vehicles 1 and 3 are in conflict. These are the two vehicles that are already waiting at the stop line and that request an opposing phase. As the intersection is unaware of the state of the neighboring intersection it foresees no conflict for the vehicles approaching the intersection from the south. The choice whether to give vehicle 1 or 3 the right of way is arbitrary as each choice will bring about the same amount of delay. In this example the intersection decides to give the right of way to vehicle 1, resulting in a delay for vehicle 3. This iteration does not end by directly acting on the choices made, but by informing nearby intersections about the intended actions. If the signal plans resulting from this iteration would be implemented this would result into four stops (although each individual intersection is only aware of one stop, so two stops in total).

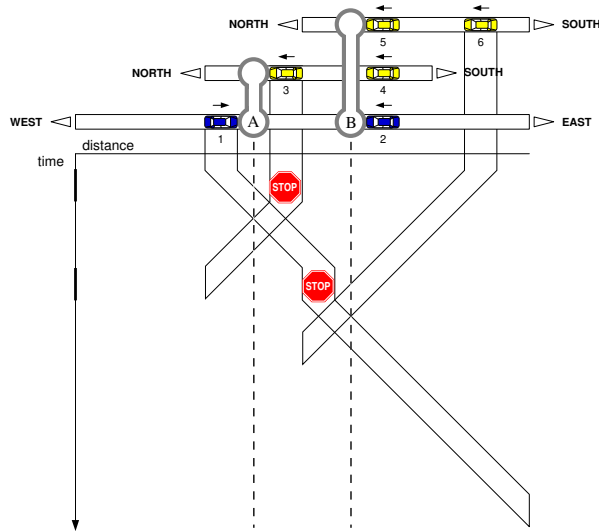
When information regarding the intentions of nearby intersections is received the next iteration starts. In this iteration the intersection observes that its choice for giving the right of way to vehicle 1 is suboptimal as the result of this choice is that the vehicle will be stopped at the next intersection. This choice results in two stops (four in total for the two intersections), whereas the decision to give the right of way to vehicle 3 leads to only one stop (two in total). The intersection agent therefore decides to give way to vehicle 3. The delay inflicted to vehicle 1 enables it to pass freely through the next junction. The intersection also observes that there is a new conflict between the vehicle 2 that originates from the nearby intersection and vehicle 4 that approaches the intersection. The choice which of these vehicles to give the right of way is arbitrary as each choice will bring about the same amount of delay. In this example the intersection decides to give way to vehicle 4. The second iteration again does not end by acting upon the choices made but instead by informing nearby intersections about its new intentions. If the signal plans resulting from this iteration would be implemented this would result into four stops (of which both intersections are now fully aware).

As soon as updated information regarding the intentions of nearby intersections becomes available the next iteration starts. This time both intersection stick to their decision to give right of way to the northbound vehicles 3 and 5. However, they also realize that there is no longer a conflict between the vehicle that originates from the nearby intersection and vehicles 4 and 6. Vehicles 4 and 6 therefore do not have to be delayed. As the intentions of each intersection were again changed the third iteration ends by informing nearby intersections about their new intentions. If the signal plans resulting from this iteration would be implemented, this would result into only two stops (of which both intersections are fully aware), compared to four stops in the previous interactions.

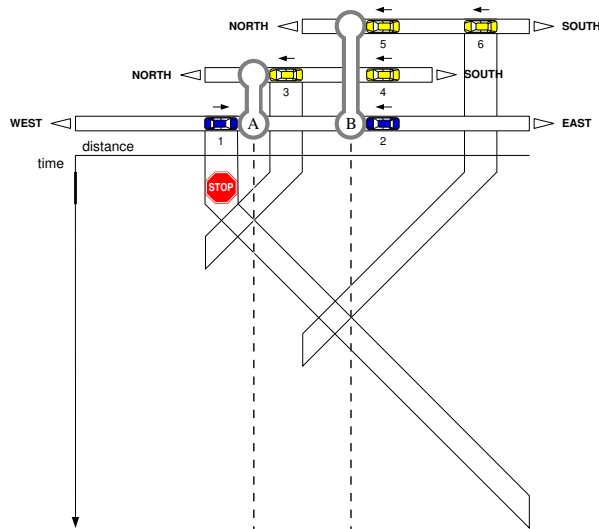
The final iteration starts as soon as updated information regarding the intentions of nearby intersections becomes available. This time the new information, for both intersections, does



(a) Conflicts on the corridor



(b) Traffic is controlled suboptimally



(c) Traffic is controlled optimally

Figure 5: Different ways in which traffic could be handled

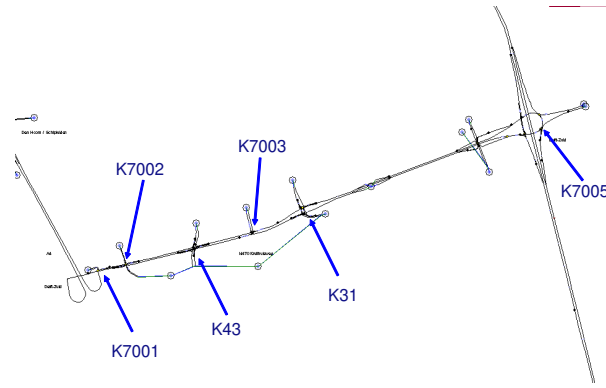


Figure 6: Geometry of the N470.

not lead to any changes in intentions. As now the process has converged, the intentions can be acted upon.

CASE STUDY

Under contract of the province of South-Holland, The Netherlands, a comparison has been made between the proposed multi-agent control approach and the traffic-actuated controllers that are currently used to control the intersections of the N470, a provincial road that connects the A13 and A4 freeway near the city of Delft.

The N470 is controlled by five traffic-actuated controllers. The N470 is depicted in Figure 6 together with the names used to denote the individual controllers. Of these, controllers K31 and K7005 control complex intersections consisting of two and four simple intersections respectively. Agent controllers are assigned to each individual simple intersection (10 in total). The behavior of the complex intersections thus emerges as a result of the coordination between the simple intersections.

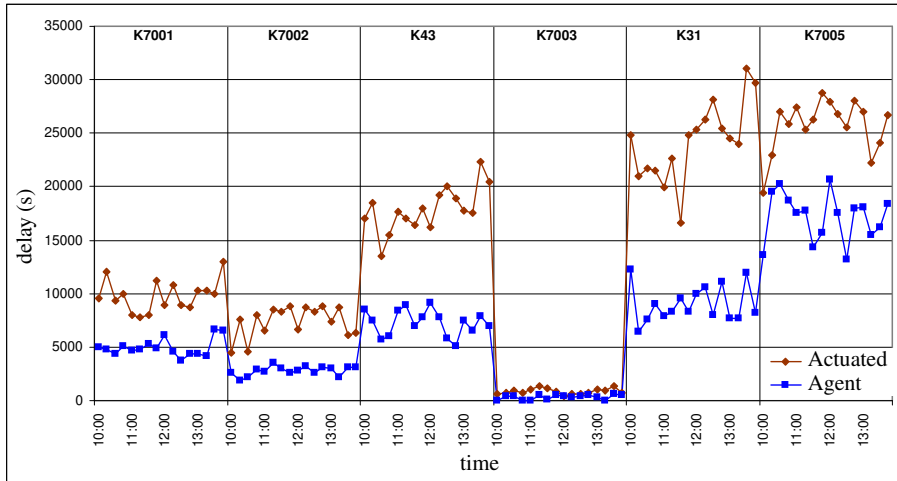
Figure 7(a) shows the added value of the multi-agent control approach with respect to the traffic-actuated control approach. The delay inflicted by the intersections when controlled by the agents in all cases is on average 50 % lower than in the reference situation.

Figure 7(b) shows the maximum queue length per quarter of an hour. Note that because the total number of vehicles waiting for a traffic signal is divided by the number of lanes, the queue length shown does not necessarily have to correspond to a whole number of vehicles. The figure shows that the travel time gained using the multi-agent control approach does not come at the expense of longer queues on approaches for which the saturation flow is less.

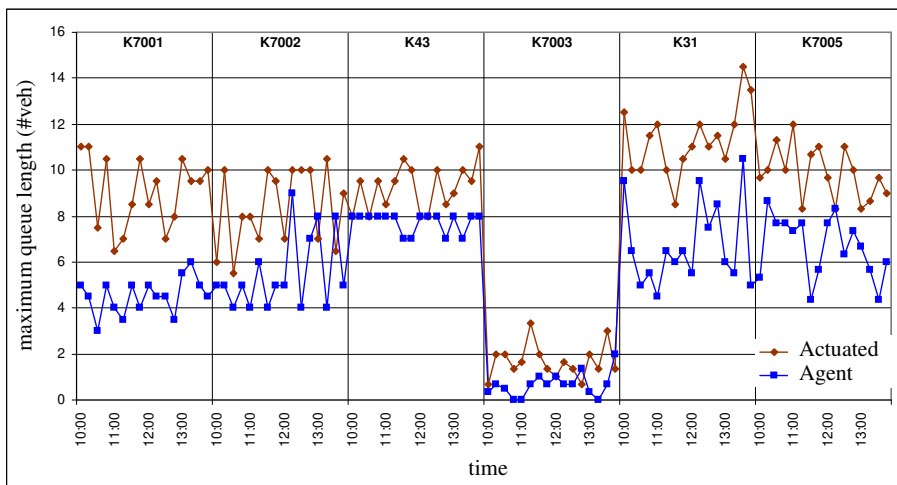
CONCLUSIONS

This paper illustrated the benefits of multi-agent coordination and defines a procedure through which coordination between traffic control instruments can be achieved. An algorithm for the look-ahead traffic-adaptive control of an intersection was furthermore presented. The algorithm integrates the currently best known dynamic programming optimization approach to look-ahead traffic-adaptive control and applies the more flexible movement-based approach to traffic signal control as opposed to the stage-based approach employed by the current state of the art in look-ahead adaptive control.

By modeling the individual traffic control instruments as intelligent agents, the actions of the individual instruments can be coordinated. The advantage of the multi-agent coordination



(a) Delay.



(b) Maximum queue length.

Figure 7: Selection of results for the N470 study.

procedure described in this paper over traditional forms of coordination, is that it is adaptive. The multi-agent coordination procedure is able to adapt to different traffic volumes and platoon ratios, and is able to create and to dissolve progression between consecutive intersections.

The developed multi-agent coordination procedure can be of considerable help in coordinating the individual autonomously functioning traffic control instruments that are deployed along the roads today. By allowing the individual instruments to coordinate their actions based on the information they receive from sensors and each other, a finer means of controlling traffic on a network can be realized. This allows the traffic management center to focus on managing the traffic network and delegate the details of coordination toward the instruments themselves.

ACKNOWLEDGMENTS

Research funded by the BSIK project “Transition to Sustainable Mobility (TRANSUMO)”, the STW VIDI project “Multi-Agent Control of Large-Scale Hybrid Systems” (DWV.6188), the Transport Research Centre Delft, the Delft Research Center Next Generation Infrastructures, and the Province of South-Holland.

REFERENCES

- [1] Hunt, P., D. Robertson, R. Bretherton, and R. Winton, *SCOOT - a traffic responsive method of coordinating signals*. TRRL Laboratory Report 1014, TRL, 1981.
- [2] Lowrie, P., The Sidney Co-ordinated Adaptive Traffic System: Principles, Methodology, and Algorithms. In *Proc. IEE Conference on Road Traffic Signaling*, 207, 1982, pp. 67–70.
- [3] van Katwijk, R., and P. van Koningsbruggen, Coordination of traffic management instruments using agent technology. *Transportation Research Part C*, vol. 10, 2002, pp. 455–471.
- [4] Mueller, E., Aspects of the History of Traffic Signals. *IEEE Transactions on Vehicular Technology*, vol. 19, no. 1, Feb. 1970, pp. 6–17.
- [5] van Katwijk, R., *Multi-Agent Look-Ahead Traffic Adaptive Control*. Ph.D. thesis, Delft University of Technology, Delft, The Netherlands, Jan. 2008.
- [6] van Katwijk, R., B. De Schutter, and J. Hellendoorn, Movement-Based Look-Ahead Traffic-Adaptive Intersection Control. In *Proc. of the International Conference on Control 2008*, Sep. 2008.
- [7] Shelby, S., Single Intersection Evaluation of Real-Time Adaptive Traffic Signal Control Algorithms. In *Proc. 84rd Annual Meeting of the Transportation Research Board (TRB'04)*, Jan. 2004.
- [8] Henry, J., J. Farges, and J. Tuffal, The PRODYN Real Time Traffic Algorithm. In *Proc. 4th IFAC/IFIP/IFORS Symposium on Control in Transportation Systems*, Baden-Baden, Germany, Apr. 1983, pp. 307–312.
- [9] Gartner, N., OPAC: A Demand-Responsive Strategy for Traffic Signal Control. *Transportation Research Record*, vol. 906, 1983, pp. 75–81.
- [10] Mauro, V., and C. Di Taranto, UTOPIA. In *Proc. 2nd IFAC-IFIP-IFORS Symposium on Traffic Control and Transportation Systems*, 1989, pp. 575–597.
- [11] Head, K., P. Mirchandani, and D. Sheppard, Hierarchical Framework for Real-Time Traffic Control. *Transportation Research Record*, vol. 1360, 1992, pp. 82–88.
- [12] Porche, I., M. Sampath, R. Sengupta, Y.-L. Chen, and S. Lafortune, A Decentralized Scheme for Real-Time Optimization of Traffic Signals. In *Proc. 1996 IEEE International Conference on Control Applications*, Sep. 1996.
- [13] Sen, S., and K. Head, Controlled Optimization of Phases at an Intersection. *Transportation Science*, vol. 3, 1997, pp. 5–17.