

Technical report 09-041

# Multi-agent control of traffic networks: Algorithm and case study\*

R.T. van Katwijk, B. De Schutter, and J. Hellendoorn

*If you want to cite this report, please use the following reference instead:*

R.T. van Katwijk, B. De Schutter, and J. Hellendoorn, “Multi-agent control of traffic networks: Algorithm and case study,” *Proceedings of the 12th International IEEE Conference on Intelligent Transportation Systems (ITSC 2009)*, St. Louis, Missouri, pp. 316–321, Oct. 2009.

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://pub.bartdeschutter.org/abs/09\\_041](https://pub.bartdeschutter.org/abs/09_041)

# Multi-Agent Control of Traffic Networks: Algorithm and Case Study

R.T. van Katwijk, B. De Schutter, and J. Hellendoorn

**Abstract**—As more and more traffic control instruments are installed, the probability increases that conflicts will arise or that coordination opportunities are lost when various traffic control measures 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 measures as intelligent agents, the actions of the individual instruments can be coordinated. By coordinating neighboring traffic control instruments, it can be prevented that they impede each other and one can realize that they function better as a whole. This paper proposes a procedure through which coordination between traffic control instruments can be achieved. We also illustrate the benefits of this multi-agent coordination approach for a simulation-based case study for the N470 arterial near Delft, The Netherlands, in which the new coordination approach is combined with a recently proposed block-based look-ahead traffic-adaptive control algorithm.

## I. INTRODUCTION

To improve traffic flow, safety, and air quality road authorities have installed many traffic control instruments, such as traffic signals, ramp metering installations, dynamic route information panels, and variable message signs (e.g., for dynamic speed control or lane closure commands). Currently, the majority of the installed traffic control instruments function fully autonomously and have 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. However, as more and more instruments are deployed, the probability increases 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 often 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 signals 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 essentially 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, which prevents the real-time coordination and timely on-line update of all the traffic control measures with optimal settings. Hence, a more distributed control framework is required. 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 distributed traffic control approach that is able to mitigate the shortcomings of a centralized approach and that can serve as a complement to the aforementioned centralized systems.

This paper is organized as follows. In Section II we recapitulate the new block-based look-ahead traffic-adaptive control algorithm for urban intersections that we have proposed in [3]. Next, we describe our multi-agent approach to the network control problem in Section III, including the procedure to coordinate the actions of the agents and an illustrative synthetic example. Subsequently, in Section IV, the results of a simulation study are presented in which we compare the new multi-agent approach with the currently implemented and well-tuned vehicle-actuated control strategy for the N470 arterial near Delft, The Netherlands.

## II. BLOCK-BASED TRAFFIC-ADAPTIVE CONTROL

In the case study of Section IV we will consider coordinated traffic signal control based on the new multi-agent approach for traffic control proposed in this paper. 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, the coordination algorithm requires a lower-level control algorithm [3] that is executed by each individual agent and that is described next. Note that we do not explicitly define a performance function as the algorithm is suited to work with various performance criteria (such as total delay, total number of stops, etc.).

Research funded by the BSIK project TRANSUMO, the STW project "Multi-Agent Control of Large-Scale Hybrid Systems", the European STREP project "Hierarchical and distributed model predictive control (HD-MPC)", contract number INFSO-ICT-223854, the European COST Action TU0702, the Transport Research Centre Delft, the Delft Research Center Next Generation Infrastructures, and the Province of South-Holland.

R.T. van Katwijk is with the Netherlands Organization for Applied Scientific Research, TNO, Delft, The Netherlands, ronald.vankatwijk@tno.nl

B. De Schutter and J. Hellendoorn are with the Delft Center for Systems and Control, Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands, b@deschutter.info, j.hellendoorn@tudelft.nl

### A. Look-Ahead Traffic-Adaptive Control

*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, which is an NP-hard problem. Hence, one has to develop efficient ways to search the decision space for (sub)optimal solutions. Efficiency in searching the decision space is characterized by the degree to which the entire tree will not have to be explored to find an optimal path. In the next subsection we will discuss a new approach for look-ahead traffic signal control that is based on blocks and that results in a much more efficient search procedure than the current approaches.

The complete low-level traffic control algorithm is based on a dynamic programming formulation similar to that of [4]. A full description of the algorithm can be found in [3].

### B. Block-Based Traffic Signal Control

Consider the intersection depicted in Figure 1 in which some possible movements of traffic (vehicles and bicycles) are indicated. A “movement” corresponds to a stream of vehicles that could get green or red, such as, e.g., movement 11 in Figure 1, 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.

Look-ahead traffic signal control then at each control step involves determining a sequence of stages to be selected from a given set of possible stages as well as their starting (and end) times. This is a mixed-integer optimization problem with the main factor of complexity being determined by the discrete optimization.

The state-of-the-art in look-ahead traffic-adaptive control (such as PROLYN [5], OPAC [6], UTOPIA-SPOT [7], RHODES [8], ALLONS-D [9], etc.) still uses a stage-based approach since it is not possible to consider all possible combinations of movements as the size of the search space grows exponentially with the number of possibilities considered. The look-ahead traffic-adaptive algorithm we propose does consider individual movements by organizing the movements into so-called blocks. In contrast to the stage-based approach where a given movement gets green (or red) for the entire




Fig. 1. Intersection used as an example




Fig. 2. Illustration of the block/movement-based approach. Top: Basic blocks; bottom: Some movements of the next block can already start before the end of the current block.

duration of the stage, we allow a movement to switch, e.g., from green to red within a block, at which time another non-conflicting movement from the next block can get green. This is illustrated in Figure 2 where the top figure shows three blocks without any premature switching yet, and where the bottom figure shows the same blocks but where some of the movements of the next block have already started before the end of the current block (indicated by the time shifts  $\tau_I$ ,  $\tau_{II}$ , and  $\tau_{III}$ ). In these figures the period of time in which a movement gets green or red is indicated by a horizontal bar that is colored accordingly. The bottom figure shows, e.g., 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. By adopting this block/movement-based approach (or block-based approach for short) we can increase the number of green combinations without increasing the size of the search space, as will be illustrated next.

Note that for a block/movement-based approach look-ahead traffic signal control then at each control step involves determining a sequence of blocks to be selected from a given set of possible blocks as well as their starting (and end) times and the times with which the green movements of the next block can be advanced. This is also a mixed-integer




Fig. 3. Number of blocks versus number of stages

optimization problem with the main factor of complexity being again determined by the discrete optimization. However, with a given number of blocks the block/movement-based approach allows to emulate a larger number of stages as will be shown next. So by adopting the block/movement-based approach we can increase the number of green combinations without increasing the size of the search space.

### C. Block-Based Versus Stage-Based Traffic Signal Control

By adopting the block/movement-based approach we can increase the number of green combinations without increasing the size of the search space as illustrated in Figure 3. This figure repeats the possible timing of green and red intervals for the movements of the intersection of Figure 1 that is presented in Figure 2 (bottom). As the block-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 thus significantly reduces the branching factor of the tree and this in its turn significantly reduces the search space without making sacrifices with regard to the quality of the solution.

## III. MULTI-AGENT TRAFFIC CONTROL APPROACH

### A. Multi-Agent Systems

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 [10] can aid in the distribution of the problem and facilitate the coordination 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 is a system comprised of several agents that together are capable to attain an objective that is difficult or impossible 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 [11]. In order to create a more intelligent traffic control system 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 subsection.

### B. Coordination Procedure

The idea of coordination between traffic control instruments has a long history [12]. 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. The 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 for 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 of 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 has to incorporate 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. While the actions of the downstream traffic control agent are considered as given, the own signal plan is the result of an iterative optimization procedure.

At first no information from neighboring control agents is available. Subsequently, the signal plan is iteratively updated to incorporate the changes at neighboring control agents as they become available. To this aim we propose the following procedure. Note that the proposed coordination approach fits

in a moving horizon or look-ahead traffic control approach in which at each time step (with an sampling interval ranging from the seconds to minute range) optimal control actions are determined over a certain prediction horizon (typically ranging from a few to several minutes).

At each time step each agent determines its current state on the basis of the information it gathers from local detectors and from information received from upstream and downstream agents. Next, 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 high-level procedure of Algorithm 1. Note that for Step 4 of the algorithm we can use the look-ahead traffic-adaptive control algorithm of Section II.

---

**Algorithm 1** High-level coordination procedure

---

```

1: loop
2:   while NOT (equilibrium reached 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

```

---

### C. 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 (A and B). The traffic demand consists of six vehicles, one eastbound approaching intersection A (vehicle 1), one westbound approaching intersection B (vehicle 2) and four northbound approaching intersection A (vehicles 3 and 4) and intersection B (vehicles 5 and 6) respectively. Blue/dark vehicles are used to indicate the vehicles that travel from the west to the east or vice versa, whereas the yellow/light vehicles indicate the vehicles that travel from the south to




Fig. 4. Geometry of the corridor used in the coordination example

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 intersection A and intersection B 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 intersection A and with vehicle 6 at intersection B. 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.

These potential conflicts can be dealt with 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 intersection A is unaware of traffic approaching the intersection B, it is unaware of the impact of its decision to the first released vehicle 1. If the intersection releases this vehicle first, the vehicle will come in conflict with vehicle 6 on intersection B. The coordination procedure proposed in Section III-B can prevent suboptimal behavior at the local intersection. The process that results from applying this coordination procedure is described next.

There is one local controller (agent) for each intersection. Initially both controllers are unaware of what the situation is at nearby intersections. As the process of intersection A mirrors that of intersection B, in the remainder only the former will be described in detail. Intersection A starts out to observe the traffic state on its approaches. The intersection

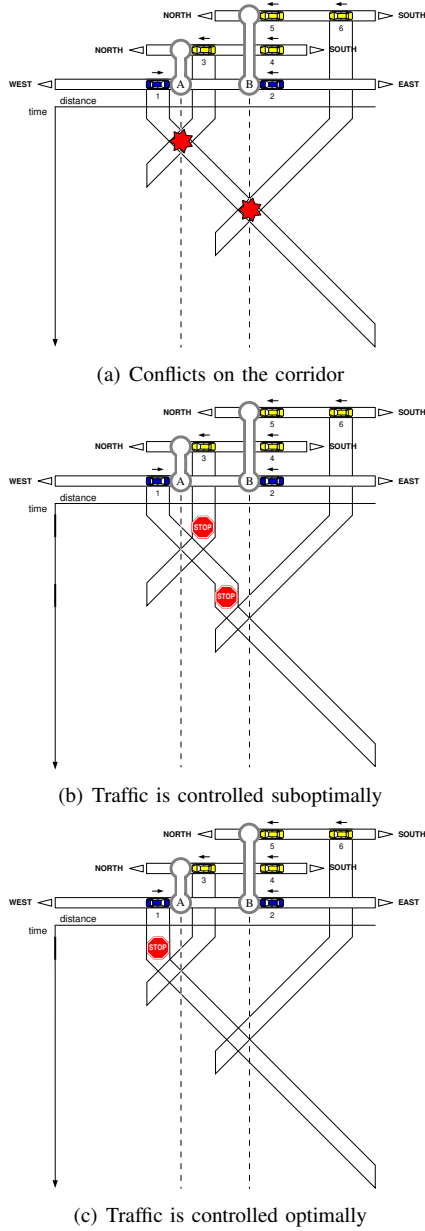


Fig. 5. Potential conflicts and different ways to handle them

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 intersection A is unaware of the state of the neighboring intersection B, it foresees no conflict there for vehicles 3 and 4. 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 actually result into four stops in total (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 intersection A 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 (B). 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. Intersection A 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. Now both intersections 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 the new intentions. If the signal plans resulting from this iteration would be implemented, only two stops would result (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 not lead to any changes in intentions. As now the process has converged, the intentions can be acted upon.

#### IV. CASE STUDY

Now we present the results of a case study in which the new multi-agent coordination approach of Section III is combined with a block-based look-ahead traffic-adaptive control algorithm (see Section II). Under contract of the province of South-Holland, The Netherlands, a comparison has been made between the proposed multi-agent control approach and the vehicle-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 Netherlands.

The N470 is controlled by five vehicle-actuated controllers. The N470 is depicted in Figure 6 together with the names (starting with “K”) used to denote the individual controllers. Of these, controllers K31 and K7005 control complex intersections consisting of two and four simple






Fig. 6. Geometry of the N470

TABLE I

AVERAGE NUMBER OF VEHICLES HANDLED PER HOUR FOR EACH OF THE INTERSECTIONS AND INTERSECTION COMPLEXES

intersection	K7001	K7002	K43	K7003	K31	K7005
volume (veh/h)	690	505	1510	1380	2490	6770

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.


We consider a typical traffic scenario for the given corridor, spanning a period from 10.00 am to 14.00 pm. Although this period is outside the rush hour, the traffic demand is still quite high during the period. The average number of vehicles handled by each of the intersections is summarized in Table I. We compare the multi-agent approach with the vehicle-actuated control that is currently controlling the intersections and that similar to the TRANSYT system [13] and that has reached a well-tuned performance over the past years. For the look-ahead algorithm we use a control sample step of 1 s and a look-ahead horizon of 2 min. As traffic simulation model AIMSUN was used.

Figure 7(a) shows the added value of the multi-agent control approach with respect to the vehicle-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.

## V. CONCLUSIONS

We have illustrated the benefits of multi-agent coordination and defined a procedure through which coordination between traffic control instruments can be achieved. By modeling the individual traffic control instruments as intelligent agents, the actions of the individual instruments



(a) Delay



(b) Maximum queue length

Fig. 7. Results for the vehicle-actuated and the new multi-agent approach

can be coordinated. The advantage of the proposed approach over traditional forms of coordination, is that the multi-agent coordination procedure is able to adapt to different traffic volumes and platoon ratios, and able to create and to dissolve progression between consecutive intersections.

## REFERENCES

- [1] P. Hunt, D. Robertson, R. Bretherton, and R. Winton, "SCOOT - a traffic responsive method of coordinating signals," TRL, TRRL Laboratory Report 1014, 1981.
- [2] P. Lowrie, "The Sidney Co-ordinated Adaptive Traffic System: Principles, methodology, and algorithms," in *Proc. IEEE Conference on Road Traffic Signaling*, ser. 207, 1982, pp. 67–70.
- [3] R. van Katwijk, "Multi-agent look-ahead traffic adaptive control," Ph.D. dissertation, Delft University of Technology, Delft, The Netherlands, Jan. 2008.
- [4] S. Sen and K. Head, "Controlled optimization of phases at an intersection," *Transportation Science*, vol. 3, pp. 5–17, 1997.
- [5] J. Henry, J. Farges, and J. Tuffal, "The PROLYN real time traffic algorithm," in *Proc. 4th IFAC/IFIP/IFORS Symposium on Control in Transportation Systems*, Baden-Baden, Germany, Apr. 1983, pp. 307–312.
- [6] N. Gartner, "OPAC: A demand-responsive strategy for traffic signal control," *Transportation Research Record*, vol. 906, pp. 75–81, 1983.
- [7] V. Mauro and C. Di Taranto, "UTOPIA," in *Proc. 2nd IFAC-IFIP-IFORS Symposium on Traffic Control and Transportation Systems*, 1989, pp. 575–597.
- [8] K. Head, P. Mirchandani, and D. Sheppard, "Hierarchical framework for real-time traffic control," *Transportation Research Record*, vol. 1360, pp. 82–88, 1992.
- [9] I. Porche, 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.
- [10] G. Weiss, Ed., *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. Cambridge, Massachusetts: MIT Press, 1999.
- [11] R. van Katwijk and P. van Koningsbruggen, "Coordination of traffic management instruments using agent technology," *Transportation Research Part C*, vol. 10, pp. 455–471, 2002.
- [12] E. Mueller, "Aspects of the history of traffic signals," *IEEE Transactions on Vehicular Technology*, vol. 19, no. 1, pp. 6–17, Feb. 1970.
- [13] D. Robertson, "TRANSYT: A traffic network study tool," Road Research Laboratory, Crowthorne, UK, Tech. Rep. LR 253, 1969.