

Technical report 08-019

Multi-agent coordination of traffic control instruments*

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 coordination of traffic control instruments,” *Proceedings of the International Conference on Infrastructure Systems 2008: Building Networks for a Brighter Future*, Rotterdam, The Netherlands, 6 pp., Nov. 2008. Paper 141.

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/08_019.html

Multi-Agent Coordination of Traffic Control Instruments

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

Abstract—As more and more traffic control instruments are installed to promote the flow in road traffic networks the probability increases that either conflicts will arise or coordination opportunities are lost when traffic control instruments are applied in the same area. 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 proposes a procedure through which coordination between traffic control instruments can be achieved.

I. INTRODUCTION

To improve traffic flow and safety road authorities have installed many traffic control instruments, such as traffic lights, ramp metering installations, dynamic route information panels and “motorway signaling” [1] (for autonomous speed control and lane closure commands). Most traffic control instruments installed have been tuned individually to attain a predetermined local objective. 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.

In order for a traffic operator to focus on the management aspects of traffic control and to ensure the timely coordination of the traffic control instruments it is desirable that certain coordination tasks are dealt with automatically whenever possible. 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.

In the literature many examples exist where the answer to the dynamic traffic control problem is sought in the form of a traffic control center that monitors the traffic network and that performs a global, or area-wide, optimization to determine new parameters for its local controllers. Much of this work [2], [3], [4] has focused on centralized, and typically predictive, control. 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.

One way to more effectively handle contingencies 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 operator can then focus on traffic management instead of on traffic control, as the latter is fully dealt with by the traffic control instruments.

This paper is organized as follows. In Section II we describe our approach to the dynamic traffic control problem. Subsequently, in Section III, this approach is illustrated for a small network. In this section it is illustrated how agents cooperate to solve a local problem in the network by taking appropriate action elsewhere in the network. In Section IV the procedure to perform the coordination among agents is introduced. This paper ends with Section V which draws some conclusions.

II. APPROACH

In [5] it is argued that multi-agent systems can aid in the distribution of the traffic control 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, which observes and acts upon an environment through sensors and actuators, whereas a multi-agent system is a system composed of several agents, collectively capable of reaching goals that are difficult to achieve by an individual agent or monolithic system.

Agent technology can make an important contribution to traffic management since the abstract concepts and ideas used in the domain of agents and multi-agent systems easily translate to the traffic management domain. 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.

The approach we choose, is to represent each infrastructure element in the network by an agent. Each agent is made responsible for the performance of the link or node in the network it represents. An infrastructure-centric approach, as opposed to a controller-centric approach, is chosen because the former abstracts more easily to the network control objectives as used by a traffic manager. The agents provide the controllers with the necessary information in order to maximize the link’s performance (in our case: minimizing the travel time spent or maximizing throughput). This information is different for each type of controller.

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, and the Delft Research Center Next Generation Infrastructures, the European STREP project “Hierarchical and distributed model predictive control (HD-MPC)”, and the European COST Action TU0702.

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, Delft, The Netherlands b@deschutter.info, j.hellendoorn@tudelft.nl

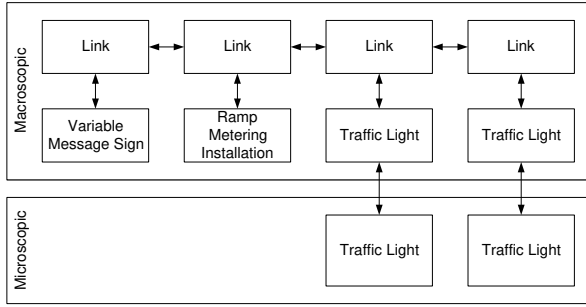


Fig. 1. Architecture of the multi-agent system




Fig. 2. Network used to illustrate the approach

The multi-agent system that is thus formed, forms a complete “shadow” network of the real network. Although each agent has a large degree of autonomy it is made susceptible to the interests of its neighbors. It therefore coordinates its actions with its neighboring agents. This coordination is obtained through direct communication and negotiation. As variable message signs and ramp metering installations can only be coordinated on a macroscopic level (i.e., on the level of flows), these controllers are not provided with any microscopic information (i.e., on the level of individual vehicles). This is different for traffic signals as traffic signals can be coordinated on the microscopic level. Figure 1 summarizes the architecture that has been chosen for the multi-agent system.

III. ILLUSTRATIVE EXAMPLE

The multi-agent traffic control concept we develop will be illustrated using the network depicted in Figure 2. This network is comprised of three separate routes (marked 1, 2, and 3 respectively). Each of these routes will get the road user from location A, which is located on the left side of the figure, to location B, which is located on the right side of the figure. Under free-flow conditions route 1 gets road user from A to B the fastest. However, road users can also choose to make use of the alternative routes 2 and 3 if they prefer. The route that the road user will ultimately choose can be influenced through means of the variable message signs that are located at each choice point in the network.

Using the network depicted in Figure 2, the benefits of coordinated control can be easily illustrated. Consider the case that an incident occurs on the preferred route (route 1). Because of this incident the preferred route is no longer able to accommodate the traffic demand. Without intervention the queue that will form upstream of the incident location will grow rapidly as more and more vehicles will join the

queue. In order to prevent that vehicles join the queue needlessly, e.g., because an alternative route would have been available to them, the agent that represents route 1 informs the upstream agent regarding the remaining capacity on the route, which is illustrated in Figure 3(a).

The agent upstream of route 1 represents the link at which traffic can choose from multiple alternative routes. This link agent is also informed by the agent representing the downstream alternative (in this case route 2), regarding any spare capacity available that can be used to accommodate the surplus of traffic. This is illustrated in Figure 3(b). If spare capacity is available the node agent can inform approaching traffic about the available alternative. This way traffic that at first wanted to make use of the preferred route (route 1), can be partly redirected to the alternative route (route 2).

When the available spare capacity is still insufficient to accommodate the surplus of traffic and route 2 as such is able to take care of only part of the problem, a solution for the remaining part must be found further upstream. This is again done by informing upstream agents regarding the remaining downstream capacity (illustrated in Figure 3(c)).

As soon as this information reaches the agent that represents the link at which traffic can again choose from alternative routes a solution is again sought at the downstream alternative (in this case route 3). This is illustrated in Figure 3(d). If the spare capacity on route 3 is sufficient to take care of the remaining surplus of traffic, then a further escalation of the problem can be prevented. Only traffic that has already made the choice for route 1 will contribute to the further growth of the queue.

Although the concept of multi-agent traffic control is easily illustrated the coordination of traffic control instruments is not a simple task. Fortunately the agents have one common goal and are thus cooperative. To achieve coordination between the agents we have developed an iterative coordination procedure, which will be described in the next section.

IV. COORDINATION PROCEDURE

In this section the procedure that is used to coordinate the agents is introduced. The coordination procedure assumes that each of the agents functions fully autonomously and tries to maximize its own, local, performance.

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 by upstream agents to downstream 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 agent can then decide whether the downstream costs outweigh its own cost, and make a decision regarding the outflow it wants to realize. If downstream costs are such that they do not outweigh its own cost, the agent can only regulate its performance




Fig. 3. Illustration of the multi-agent traffic control concept

by manipulating the inflow. This is done by incorporating the downstream outflow costs into the inflow cost the agent attributes to upstream agents.

As not all links are equipped with traffic control instruments, agents more often than not have no other means of control than to increase the cost for inflowing traffic for upstream links. Problems can therefore often not be dealt with close to the source of the problem, but can only be dealt with by actuators elsewhere in the network.

The procedure described is used to coordinate the actions of the traffic control instruments in the network on both a macroscopic and the microscopic level and is summarized using the high-level procedure of Algorithm 1.

In steps 3–4 of the procedure agents try to optimize their local performance. Coordination with neighboring agents takes place in steps 5–8 for downstream agents and in steps 9–10 for upstream agents. The coordination procedure is re-

Algorithm 1 High-level coordination procedure

```

1: loop
2:   while NOT (converged 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

```

peated until the agents no longer wish to update their control settings. The coordination procedure is further elaborated in Section IV-A and Section IV-B for coordination on the microscopic and macroscopic level respectively.

A. COORDINATION ON A MICROSCOPIC LEVEL

The idea of coordination between traffic control instruments is very old [6]. 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 has to be developed [7], [8], [9], [10] 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.

To come up with a progression scheme that is able to flexibly adapt to changes in volumes and platoon ratio a traffic signal 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. The benefits and the dynamics of the microscopic coordination procedure 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 (1), one westbound approaching the eastern intersection (2) and four northbound approaching the western (3,4) and eastern (5,6) intersections respectively. Blue/dark vehicles are used to denote the vehicles that travel from the west to the east or vice versa whereas the yellow/light vehicles denote the vehicles that




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

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 it releases this vehicle first, it will come in conflict with vehicle 6 on the eastern intersection. The coordination procedure can prevent suboptimal behavior at the local intersection. The process which 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

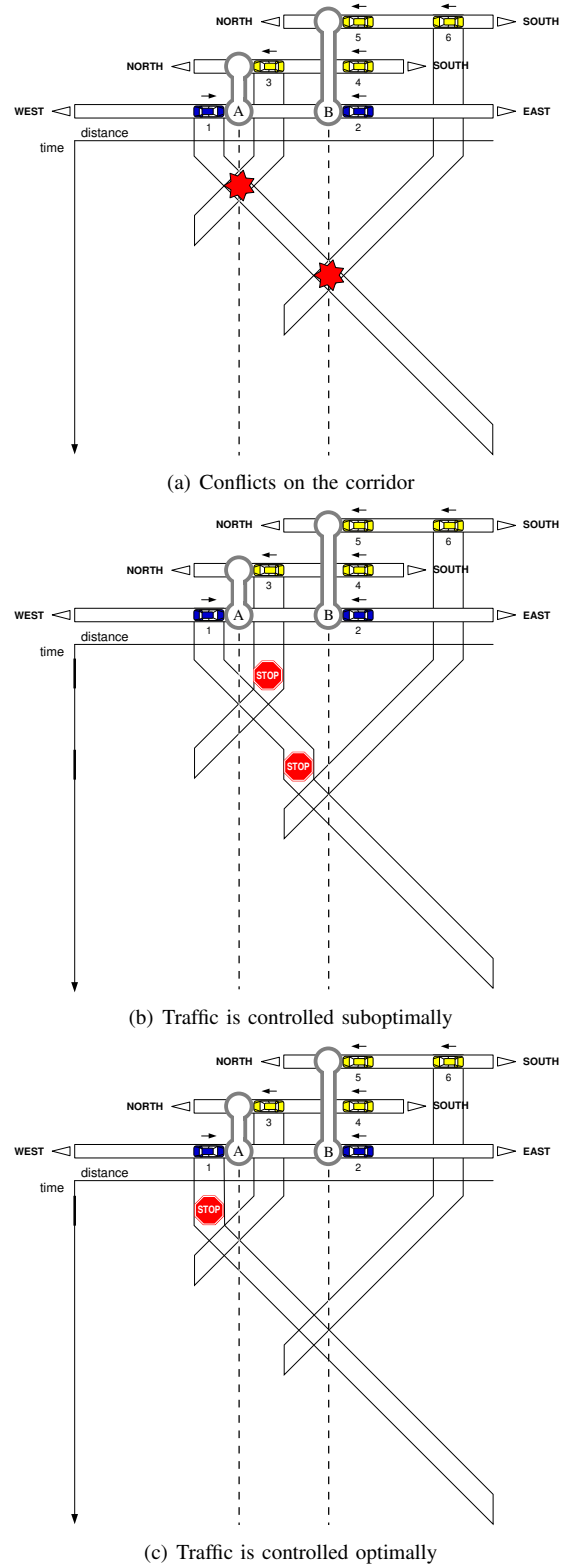


Fig. 5. Different ways in which traffic could be handled

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 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 delay for vehicle 3. This iteration ends not by directly acting on the choices made, but by informing nearby intersections about the actions intended. If the signal plans resulting from this iteration would be implemented this would result into four stops (although both intersections combined are aware of only two).

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). It therefore decides to give way to vehicle 3. The delay incurred 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 the vehicle inbound from the nearby intersection. The second iteration again ends by not directly acting on 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 nearby intersection's intentions comes available the next iteration starts. This time both intersection stick by 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).

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. The intentions can therefore be acted upon as the process can be said to have converged.




Fig. 6. Geometry of the corridor used in the macroscopic coordination example

B. COORDINATION ON A MACROSCOPIC LEVEL

Many traffic-adaptive systems [2], [11], [3], [4] allow the protection of vulnerable areas in the network by holding traffic further upstream. This functionality is called “gating”. The form of gating applied in the state of the art in urban traffic control (with the exception of TUC, which claims [4] inherent gating behavior) however relies on the operator to determine a set of fixed rules used to restrict the inflow of traffic in a certain area whenever traffic counts in the protected area exceed a certain level. The “gating” applied by the macroscopic coordination procedure defined in this section works dynamically.

The agents representing the infrastructure elements of the network continuously monitor whether their part of the network is either under-utilized or over-utilized and try to direct traffic to or steer traffic away from these parts of the network. The coordination procedure developed employs two types of constraints: hard constraints enforcing that the volume of traffic entering the vulnerable area does not exceed the volume the infrastructure is able to handle, and soft constraints used to tempt agents further upstream in the network to steer traffic away from the vulnerable area, such as to alleviate the stress put on downstream agents that have already started gating. The benefits and the dynamics of the coordination procedure at a macroscopic level can be understood by considering the following example.

Figure 6 shows a corridor consisting of two intersections. This example will focus on traffic that travels from the left side to the right side of the picture. Relevant agents representing the links of the urban corridor used in the example are labeled 1–7.

Figures 7(a)–7(d) are used to illustrate the consecutive steps taken by the agents to reduce the volume of the flow originating from the far left of the corridor as a result of a capacity drop on the far right of the corridor. Assume that the state of the network starts out as illustrated in Figure 7(a). In this figure each vertical bar denotes the outflow volume requested by the upstream agents as well as the inflow granted to the upstream agents. Figure 7(a) shows that downstream agents 1 and 3 have just barely been able to fulfill the outflow requested by their upstream agents 2–4 and 5–7 respectively, which indicates that the downstream

capacities are fully utilized. Capacity levels are however sufficient to accommodate the demand.

In Figure 7(b) the outflow capacity of agent 1 has been reduced (i.e., as a consequence of the increased inflow of traffic from the minor directions on the downstream junction). Agent 1 is therefore no longer able to fulfill the outflow demand from its upstream agents 2–4.

Consequently agent 3 is also no longer able to accommodate the outflow requested by its upstream agents 5–7. This is illustrated in Figure 7(c). Whereas the downstream capacity for agents 5–7 first sufficed to accommodate the total outflow requested by these agents, agent 3 is now no longer able to process the requested demand.

In Figure 7(d) agent 3 no longer grants the outflow requested by agents 5–7 and as a consequence a queue starts building up at the link represented by agent 6.

The macroscopic coordination procedure ensures that the amount of traffic released toward the urban area is buffered outside the network if demand levels exceed the level of demand the network can handle. As traffic is buffered at the gates of the urban area (formed by agent 6) the build-up of queues of idling vehicles in the vulnerable urban area is prevented.

V. CONCLUSIONS

This paper has 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 can be coordinated. The advantage of the multi-agent coordination procedure described in this paper over traditional forms of coordination, is that it is able to adapt to different traffic volumes and platoon ratios, to create and to dissolve progression between consecutive intersections, and to restrict inflow to vulnerable parts of the network when required.

The developed multi-agent coordination procedure can be of considerable help in coordinating the individual autonomously functioning traffic management measures 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 operator to focus on managing the traffic network and to delegate the details of coordination toward the instruments themselves.

REFERENCES

- [1] J. Klijnhout, "Motorway control and signalling: The test of time," *Traffic Engineering and Control*, vol. 25, no. 4, April 1984.
- [2] P. Martin, J. Perrin, B. Chilukuri, C. Jhaveri, and Y. Feng, "Adaptive signal control ii," University of Utah, Tech. Rep. UT-03.28, 2003.
- [3] F. Busch and G. Kruse, "MOTION for SITRAFFIC - a modern approach to urban traffic control," in *Proc. Intelligent Transportation Systems*, Oakland, CA, USA, August 2001, pp. 61–64.
- [4] C. Bielefeldt, H. Condie, E. Kosmatopoulos, A. Richards, M. McDonald, J. Mück, and A. Hanitzsch, "SMART NETS final report," Napier University, Tech. Rep., 2004.
- [5] 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.
- [6] E. Mueller, "Aspects of the history of traffic signals," *IEEE Transactions on Vehicular Technology*, vol. 19, no. 1, pp. 6–17, February 1970.
- [7] J. Little, M. Kelson, and N. Gartner, "MAXBAND: A program for setting signals on arteries and triangular networks," *Transportation Research Record*, vol. 796, pp. 40–46, 1981.
- [8] N. Gartner, S. Assmann, F. Lasaga, and D. Hout, "A multi-band approach to arterial traffic signal optimization," *Transportation Research Part B*, vol. 25, pp. 55–74, 1991.
- [9] S. Venglar, P. Koonce, and T. Urbanik II, *PASSER V*, Texas Transportation Institute, 2000.
- [10] D. Hale, *Traffic Network Study Tool, TRANSYT-7F, United States Version*, McTrans Center, University of Florida, Gainesville, Florida 32611-6585, January 2006.
- [11] A. Al-Mudhaffar, "Impacts of traffic signal control strategies," Ph.D. dissertation, Royal Institute of Technology, Stockholm, Sweden, 2006.




Fig. 7. An illustration of macroscopic coordination