Available online at www.sciencedirect.com

ScienceDirect

Transportation Research Procedia 62 (2022) 236–244

Transportation Research Procedia

www.elsevier.com/locate/procedia

24th Euro Working Group on Transportation Meeting, EWGT 2021, 8-10 September 2021, Aveiro, Portugal

Multiagent Meta-level Control for Adaptive Traffic Systems: A Case Study

Yaroslava Shynkar[a], Anita Raja[a,*], Ana L. C. Bazzan[b], Marin Marinov[a]

[a]Hunter College, City University Of New York, NY, NY, USA
[b]Instituto de Informática, UFRGS, Caixa Postal 15064, 91.501-970 Porto Alegre, Brazil

Abstract

As cities across the globe continue to grow, traffic congestion has become globally ubiquitous with great economic and environmental costs associated with it. The increasing prevalence of self-driving vehicles creates an opportunity to build smart, responsive traffic infrastructure of the future. Such an infrastructure consisting of connected and autonomous vehicles and smart traffic lights would have the potential to cope with congestion, weather phenomena and accidents, while maintaining safety and ensuring privacy of information. This paper introduces an approach to address the challenge of dynamically adjusting traffic to the changes in the environment. We argue that multiagent meta-level control (MMLC) is an effective way to non-myopically determine how and when this adaptation should be done. The approach highlights the role of dynamic meta-reasoning in a platooning scenario, in which collaboration contributes to improved travel time for vehicles in the network as well as a positive environmental impact as related to fuel consumption and emissions. Specifically, for the case study described in the paper, our MMLC-based approach leads to approximately 44% decrease in travel time, 7% increase in average speed, a 32% decrease in fuel consumption and a 35% drop in emissions. We also see performance advantages for a scaled-up mixed traffic simulation environment.

Peer-review under responsibility of the scientific committee of the 24th Euro Working Group on Transportation Meeting (EWGT 2021)

Keywords: Intelligent Transportation Systems; Smart Cities; Multiagent Systems; Meta-Level Reasoning

1. Introduction

Global urbanization over the past decades has made traffic congestion ubiquitous resulting in economic and environmental costs in the hundreds of billions per year. There is a great need to leverage the increase in connected, autonomous vehicles and smart traffic lights while planning the mixed traffic infrastructure of the future. This infras-

* Anita Raja. Tel.: +1-212-650-5031.
E-mail address: anita.raja@hunter.cuny.edu (Anita Raja).

E-mail address: anita.raja@hunter.cuny.edu

2352-1465
Peer-review under responsibility of the scientific committee of the 24th Euro Working Group on Transportation Meeting (EWGT 2021)
10.1016/j.trpro.2022.02.030

tructure should be able to adapt to the dynamics and constraints of an open environment. We propose a multiagent system-based approach that can effectively determine how and when this adaptation should be done. We highlight the role of dynamic meta-reasoning in a multiagent system with respect to connected and autonomous planning and coordination in the context of limited resources. A key component in this vision involves leveraging the potential deployment of connected and autonomous vehicles (CAVs). These will change the way vehicles interact amongst themselves, as well as the way the road infrastructure is organized. CAVs have a potential to mitigate congestion, increase the throughput, facilitate the sharing of vehicles as well as autonomous mobility-on-demand(AMoD), and increase safety while maintaining privacy for all users of the transportation network.

We argue that an important aspect of the introduction of CAVs (and the necessary connected infrastructure) is that it facilitates the reorganization of the whole traffic system, so that (sub)parts can cope better with the complex interactions that occur within the system. In this paper, apart from the physical urban network that is ubiquitous in the traditional view, we also consider several other components that are networked to form a MAS. These components include a network composed of the CAVs which is managed by a set of managers of CAVs (CAVMNGRs) and capable of platoon formation. A platoon is a set of autonomous vehicles that is treated as a single unit. It resembles a train structure in which the train cars form a linked entity (Sala and Soriguera, 2020). Recent works argue that platooning contributes to increased efficiency of the traffic system (Heinovski and Dressler, 2018). Further, in keeping with the current view of connected infrastructure, we envision networks between intersection managers (IMNGRs), where intersections include traffic signals, as well as networks between link managers (LMNGRs) that manage links. Links are important parts of the network, that have received less attention. Other actors that may be networked and managed are mobility-on-demand and bike managers. The managers of the various networks could also be connected to form a heterogeneous transportation network.

So far little has been investigated regarding how to coordinate the individual objectives of such actors, with few exceptions. One is the body of research dealing with selfish routing in the context of traffic assignment. Selfish routing (Koutsoupias and Papadimitriou, 1999; Roughgarden and Tardos, 2002), where individual agents make uncoordinated greedy routing decisions, is known not to produce a socially desirable outcome in transport and communication networks. Another is the work on routing in the context of logistics. While important, these research threads mostly do not consider other actors in an integrated way.

We propose using multiagent meta-level control (MMLC) (Cheng et al., 2013; Gordon et al., 2008; Yamane and Ishida, 2006) as a necessary feature for complex agents to operate, communicate and coordinate in these heavily networked future transportation environments. These environments will inherently be more open and dynamic, with more players and disrupting technologies. MMLC is the process that facilitates an agent network (e.g., CAVMNGR or IMNGR network or network of networks) to have a decentralized meta-level multiagent policy, where the progression of what deliberations the agents should do, and when, is choreographed carefully and includes branches to account for what could happen as deliberation plays out. For example, suppose multiple intersections belonging to one IMNGR have been shut down due to a water main break. The IMNGR agent decides to devote significant resources to coordinate with its two neighboring IMNGRs and reorganize the routing of traffic. However, if those neighbors have allocated little resources for this coordination, the IMNGR agent then would waste resources trying to coordinate with an unwilling partner. It is thus important to coordinate the partners at the meta-level to ensure network reorganization.

While our long-term goal is to apply distributed MMLC to the traffic reorganization problem, for the purposes of this paper we investigate centralized MMLC and discuss its advantages and disadvantages. Our proposed hypothesis is that in dynamic traffic contexts, where flow is interrupted (due to the weather conditions, volume of vehicles change, accidents, etc.), MMLC and platooning allow for efficient re-organization of the various elements within the traffic system that lead to performance advantages with respect to speed, fuel usage and emissions. We will first describe related work in connected vehicles and future of transportation networks. We then describe our approach that uses MMLC and platooning and provide an evaluation in the context of a motivating example as well as for a scaled-up mixed traffic network. We finally discuss the conclusions and future work.

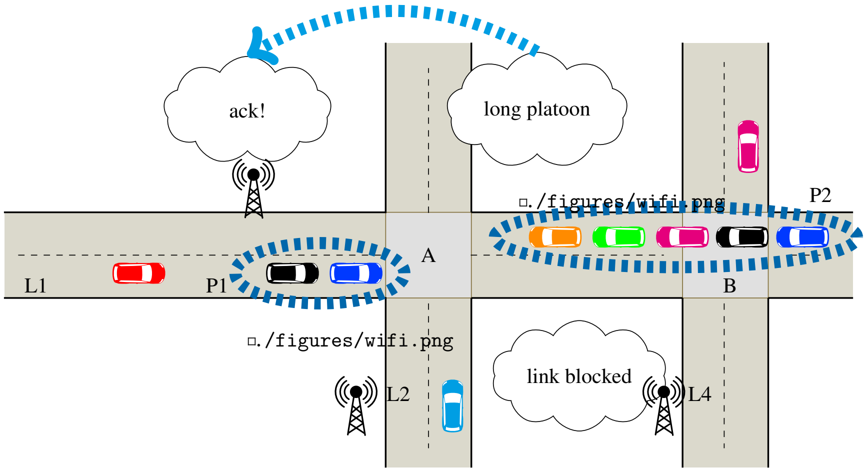

Fig. 1. A snapshot of interactions among CAV and link managers in our system

2. Related Work

The literature on the various aspects related to future of traffic and transportation systems is vast. While a systematic survey is thus outside of scope, we briefly shed light on works that relate to multiagent systems (MAS). We show that, within the MAS literature, there is a gap motivating the need for a framework capable of meta-level decision making.

Reservation-based control of intersections were addressed by (Dresner and Stone, 2004; Vasirani and Ossowski, 2009) (and their extensions). In (Au et al., 2011), a management policy is proposed, that puts the request messages on hold and then processes several requests at once, in order to avoid starvation at secondary roads. With respect to the use of navigation and traffic apps, in (Hasan et al., 2016), the authors have investigated traffic coordination algorithms for connected vehicles. Klügl and Bazzan (2021) propose the use of an app for allowing agents who share origins and/or destinations to communicate their rewards (travel time) in order to gain more information and accelerate the learning process. However, none of these works deal with decision-making at the meta-level.

CAVs are key components in our approach. In the interests of space, we refer the reader to Mahmassani's overview(Mahmassani, 2016) on the implications of CAVs for traffic and urban systems, as well as for a discussion on related works. This focuses on other important issues such as sensitivity related to market penetration of CAVs, mixed traffic performance, operational control design, and performance of the overall traffic system, among others. It also introduces an approach that focuses on flow and operational aspects of CAVs. In this sense, it has a different purpose than ours. We focus on the interactions among the various heterogeneous components of the infrastructure that underly the CAVs, besides the organization of the CAVs themselves (as, e.g., platooning). Recent surveys point out that interaction and organizational issues are not receiving much attention in current work. For instance, Khayatian et al. (2020) focus mainly on intersection management, communications, and safety and robustness issues, not on reorganization of the network via decision-making at vehicle, platoon, intersection or other aggregated levels.

3. Motivating Example and Approach

We describe a hierarchical MAS approach for future transportation networks. It is characterized by a supervisory framework for organizational control to coordinate the vehicular network while coping with congestion, weather phenomenon and accidents. Each acting agent consists of two levels of reasoning (Cheng et al., 2013). The higher level of control called meta-level control, can reason about alternate methods for deliberation (computation) and determines parameters for whether, when and how to (re)organize the network and coordinate with other agents. Specifically, the meta-level control reasoning component is addressing questions such as: whether and when to reorganize the network. The second level of control is the deliberative level which is more computationally intensive and includes algorithms for coordination, platooning, routing, signalling, etc. It determines the actions to be taken at the ground or domain level by the agent's effectors.

Figure 1 describes a portion of a road network with a set of CAVs traveling on it. It is assumed that each CAV has a valid route from an origin to a destination. The CAVMNGRs network consists of multiple $CAVMNGR_i$ agents, that cover non-overlapping areas. Two $CAVMNGR_i's$ are neighbors if they share boundaries of coverage. For the sake of illustration, consider that at a given time a platoon of two vehicles (platoon P1 in the figure) is heading to intersection A (covered by $CAVMNGR_A$) and plans to turn south at intersection B. In the meantime, a platoon of five vehicles (platoon P2 in the figure) is passing through intersection B (covered by $CAVMNGR_B$) planning to head west. Suppose $LMNGR_4$ informs $CAVMNGR_4$ that link L4 is blocked south-bound. Given this information about L4, the metalevel reasoning questions to be handled by $CAVMNGR_B$ are whether the platoon P1 should continue as is or be rerouted as a whole unit, or should it be disbanded and have all its vehicles redirected on a case-by-case basis using different routes to mitigate the congestion. For instance, the MMLC and its associated deliberative control could determine that platoon P1 should turn south at intersection A instead of proceeding east as planned. This is determined by comparing the utility of alternative routes and picking the one that has the greatest expected utility.

Also the *LMNGR* to the top left is made aware by the meta-level communication that platoon P2 is heading its way. It will coordinate with *IMNGR* to its west about this platoon. The *LMNGR* and *IMNGR* will then coordinate to determine whether platoon P2 should disband or stay together based on its expected utility values and then update the meta-level control with their decision. The deliberative-level reasoning involves applying the appropriate platooning strategy. While P2 is not directly affected by the closure of L4, there are secondary effects if the traffic interruption is not addressed. Without MMLC all vehicles intending to turn into lane 4 would get stuck and that would cause traffic congestion at intersection B. This would affect P2's journey through intersection B.

In this paper, we introduce meta-level control and platooning in CAVs to implement the above described meta-level reorganization that facilitates dynamic rerouting to handle traffic interruption events. We study their effect on a number of performance metrics of the traffic network. In our experiments, we show that platoons and non-myopic MMLC decisions to handle dynamic changes in the environment such as street shutdowns or accidents lead to a performance advantage.

Traffic Interruption Events: We consider three types of traffic interruption (TIE) events: weather, volume-change and accidents. Each event has three weights associated with it to reflect its priority: *Traffic Interruption Event (TIE) weight* (TIE_w) which indicates the priority of one event type over another; *Deadline weight* (D_w), which indicates the urgency of the TIE deadline; and *Severity weight* (S_w), which indicates the impact of the TIE on the network performance. For instance, a car accident that blocks an intersection is more severe than one that closes a small part of one lane. S_w can take on the value "low", "medium" or "high". In this paper, our MMLC approach determines how the TIEs should be handled by leveraging these weights.

Platooning: Cooperative Adaptive Cruise Control (Heinovski and Dressler, 2018) is a set of strategies (protocols) that guarantee that by forming a platoon vehicles can minimize their inter-vehicle distance and hence their speed differences. Cooperative Adaptive Cruise Control (CACC) (Wang et al., 2019) is an extension of ACC. In our work it is assumed that only CAVs can form a platoon since they are capable of sharing data (through V2X) and have the capacities that are essential for a platoon formation. In the simulation environment, if two vehicles have a similar route (as determined by the CAVMNGRs at every edge), then each vehicle is enabled for platoon formation and the merging maneuver is performed.

Multiagent Meta-level Control: The deliberative-level actions that Multiagent Meta-level Control (MMLC) has to choose from are: (a) reroute the vehicles; (b) reroute platoons; and (c) disband the platoon and reroute the vehicles that belonged to it. Each MMLC state captures information about the current time, contents of event priority queue (EPQ), an event trigger list, average travel time so far, and the MMLC decision from the previous time step. The EPQ is the ordering of events (TIEs) that could interrupt traffic flow including weather, volume change and accident events. MMLC determines which events need to be addressed first. We assume the TIE_w is a three tuple for prioritizing weather, volume-change and accident respectively. The default $TIE_w = [0.2, 0.5, 1.0]$ indicates accidents have highest priority with weight 1.0 followed by volume change with weight 0.5 and finally weather with weight 0.2 ; the S_w is a tuple for low, medium and high levels of severity respectively with default values for $S_w = [0.2, 0.4, 0.8]$; the DL_w is a tuple for relaxed and tight deadlines with default values for $DL_w = [0.25, 0.75]$. Events are ordered in the EPQ using the following intuitive weighting formula: $TIE_w * (DL_w + S_w)$.

We introduce a deterministic MMLC algorithm that uses the above weighting heuristic, which is manually set to enforce the following intuitive rules: accidents are given higher priority (higher TIE weights) over weather and volume change; if two events in the queue are of the same type, the one with the higher severity is prioritized; if two events in the queue have the same severity, then one with a tighter deadline is prioritized; if MMLC is not activated, instead of being rerouted, the vehicles simply form a queue at the junction next to the blocked edge until the event clears; if MMLC is activated, then the possible action choices based on system state are (a) reroute the vehicles; (b) reroute platoons and (c) disband the platoons and reroute the vehicles that belonged to the platoon.

In this MMLC algorithm, the EPQ accounts for long-term effectiveness of reorganization of the network. Instead of taking myopic/greedy actions to handle incoming events, contextual information, including information about other events in the queue, is used in making the action choice. This enables MMLC decisions to be non-myopic.

4. Empirical Evaluation

Control Strategies: We evaluate the effectiveness of deliberative-level platooning and the higher level MMLC in handling the consequences of traffic interruption events by comparing four control strategies in a variety of environments. The four strategies are: **Strategy 1 (S1)** where vehicles in the network have no option for platooning or MMLC and the vehicles individually respond to TIEs as they occur with no planning; **Strategy 2 (S2)** where vehicles in the network can perform maneuvers such as merge, disband, or leave a platoon at the deliberative level, but have no MMLC to handle the dynamic changes; **Strategy 3 (S3)** where MMLC is activated at the meta-level for vehicles but there is no deliberative control option for platoon formation; **Strategy 4 (S4)** where both MMLC and deliberative-level platooning are available as control options for CAVs in the network.

Microscopic Simulator and Evaluation Metrics: All experiments described in this paper use SUMO (Simulation of Urban MObility) (Lopez et al., 2018), an open source, microscopic traffic simulator. For our first set of experiments, the example depicted in Figure 1 was modelled in SUMO. We simulate and track 10 CAVs with routes that were computed with respect to their origins and destinations as well as their departure times. For our simulation, we compared the efficiency of MMLC and platooning using two car-following models available in SUMO (namely IDM and Krauss (Treiber, 2013)), taking into account their specific parameters. The Krauss model (also known as the stochastic version of Gipps' model) considers realistic parameters that capture close-to-accurate microscopic behaviour of a vehicle in highway traffic. It is collision-free since accidents are minimized due to the safe speed parameter. The Intelligent Driver Model (IDM) is another popular model that generates better macroscopic traffic network data for evaluation purposes. In comparison to the Krauss model, IDM computes a more realistic travel time for a vehicle to cross an intersection. While it does not capture the leader vehicle's speed, it keeps track of the gap between the vehicles. IDM is more suitable for CAVs because it is able to handle small gaps between vehicles and also produces realistic acceleration results. This enables it to capture more accidents than Krauss and track the fuel consumption and CO_2 emissions. IDM also makes the SUMO emissions model more suitable for microscopic situations.

When using platoon formation strategies, the default settings of our models are modified to simulate the behaviour of merging of autonomous vehicles into a platoon, i.e., CAVs in a platoon have the minimum gap set to zero, minimum time headway of 0.05, the speed is the current speed of the platoon leader vehicle. When two vehicles on the same lane perform a merge into a platoon maneuver, the routing is modeled such that the follower always has to accelerate in order to form the platoon with the leader.

Our *first goal* is to determine the usefulness of control strategies that use MMLC and platooning independently, and also when they are combined. We measure the average speed and average time to destination for the vehicles in the network. Our *second goal* is to evaluate the effect of the strategies on vehicle emissions. The SUMO emissions model captures specific pollutants emitted by the vehicle including the amounts of Carbonmonoxide (CO), a known toxic substance; Carbondioxide (CO_2), a greenhouse gas; Hydrocarbons (HC) that lead to smog and ground-level ozone; particulate matter (PM_x) that are known to cause cancer; and Nitrousoxide (NO_x) that is emitted in milligrams by the vehicles on a lane/edge, creates smog and contributes to acid rains. The emission values presented are the average for the network (averaged based on data per lane).

Mixed traffic networks consist of CAVs and manually-driven non-autonomous vehicles. Therefore, our *third goal* is to analyze the changes in performance as a function of the ratio of CAVs equipped with MMLC and Platooning in a mixed traffic network. For this purpose, we employ various metrics as discussed in the next section.

Empirical Results In these first set of experiments, we explore the performance (Table 1) and emissions (Table 2) implications of the four above-mentioned strategies when lane L4 in the simple motivating example (Figure 1) is shut down due to an accident. When using S1 and S2, where there is no MMLC and thus a lack of coordination among the agents about the TIE, it is observed that there is traffic backup in the lanes between intersection A and B and the vehicles simply wait for L4 to be open again. However with MMLC (S3 and S4), the ability to coordinate leads to the routes involving the blocked link L4 to be replaced by the MMLC recommended routes that address the lane closure and limit the traffic congestion.

	AvSpd(m/s)		AvTT(s)		Fuel Usage (ml/s)	
Col. # →	1	2	3	4	5	6
Sce. # ↓	Krauss	IDM	Krauss	IDM	Krauss	IDM
S1	11.64	11.68	35.05	34.81	69.42	59.56
S2	11.81	11.8	34.66	34.60	70.2	59.38
S3	12.73	12.73	19.25	19.2	46.71	37.58
S4	12.91	12.88	18.94	18.97	47.16	36.91

Table 1. Comparison of average vehicle speed, average travel time and fuel consumption under various control strategies (S1–S4) using the Krauss and IDM car-following models in the example described in Figure 1. The strategies are S1: No-MMLC and No-Platooning; S2: No-MMLC with Platooning; S3: MMLC only with No-Platooning; S4: MMLC with Platooning.

	CO(mg/s)		CO_2(mg/s)		NO_x (mg/s)		HC(mg/s)		PM_x(mg/s)	
Col. # →	1	2	3	4	5	6	7	8	9	10
Sce. # ↓	Krauss	IDM	Krauss	IDM	Krauss	IDM	Krauss	IDM	Krauss	IDM
S1	4813.04	4312.93	161487.88	138550.53	67.77	56.55	25.69	22.58	3.18	2.53
S2	4724.77	4334.90	163306.23	138135.75	67.06	56.49	25.11	22.67	3.04	2.54
S3	1415.21	927.74	108670.57	87413.68	43.41	32.74	8.98	5.93	1.83	1.2
S4	1233.62	900.9	109704.47	85876.93	42.31	32.21	7.96	5.78	1.67	1.19

Table 2. Comparison of emissions averaged over the network under various control strategies (S1–S4) using the Krauss and IDM car following models in the example described in Figure 1. The strategies are S1: No-MMLC and No-Platooning; S2: No-MMLC with Platooning; S3: MMLC only with No-Platooning; S4: MMLC with Platooning.

Columns 1 through 4 in Table 1 show the comparison of the average speed and average travel time for each of the four strategies and two car following models: Krauss and IDM. For instance, in column 1 we see that using the Krauss model, the average speed using strategy S3 is at least 9.36% and 10.91% higher or more than S1 and S2 respectively. Overall, we note that for both the Krauss and IDM models, strategies 3 (S3) and 4 (S4) have better performance (at least 7.79% increase in average speed and 44.46% decrease in average travel time) than strategies 1 (S1) and 2 (S2). These results show that using MMLC (S3 and S4) has a performance advantage compared to not using MMLC (S1 and S2) indicating that the ability of MMLC to adapt to changes and dynamically reroute vehicles leads to less congestion even in the simple example in Figure 1. We observe that platooning (S2 and S4) has a performance advantage (about 1-2%) compared to not using platooning (S1 and S3) respectively on almost all measures. S4 has slightly better performance than S3 even though MMLC is available in both because S4 also employs platooning. This indicates that the meta-level recommendation to reroute the platoon P1 at intersection A due to the closure of intersection B leads to the best performance with regards to average speed and lower average travel time in the simple example network.

We also observe in Columns 5 and 6 in Table 1 that fuel usage decreases significantly (by about 32% at least) with the use of MMLC and platooning regardless of car following model. This is because of the reduced aerodynamic drag since the minimum gap between the vehicles is very small (it is set to 0) under platoon formation settings. Finally, we note that it is computationally less expensive to recompute a route for all vehicles within a platoon using S2 and S4 (30 invocations of the Dijkstra rerouting algorithm within SUMO) than to recompute routes for each individual vehicles as in S1 and S3 (80 invocations of the rerouting algorithm).

Table 2 shows that, in general, the average emissions of CO, NO_x, HC and PM_x for the network reduce at least by about 70%, 35%, 64% and 39% respectively with the use of platooning and MMLC in S4 compared to S1 and S2. This is because with MMLC, fewer vehicles are stopped or slowed down when there is an unexpected interruption of traffic flow. Also with platooning leading to increased average speed, most emission gases are reduced. Our results for CO_2 in the Krauss model show a slight increase in emissions for S4 over S3. We argue that this captures the disadvantage of exclusively using the acceleration feature during platooning in an environment where all CAVs have their settings changed during the platooning, as it is the case in using S4 on the motivating example environment. The results for IDM model show that platooning can reduce the fuel consumption as well as the CO_2 emissions. This is due to the fact that IDM has parameters similar to the settings of vehicles under the platoon formation strategies.

	AvSpd(m/s)		AvTT(s)		Fuel Usage (ml/s)	
Col. # →	1	2	3	4	5	6
CAVs % ↓	Krauss	IDM	Krauss	IDM	Krauss	IDM
0	12.14	11.77	40.58	43.36	3038.06	2379.53
0.25	12.38	12.03	39.66	41.08	2813.63	2178.97
0.5	12.53	12.04	39.21	41.29	2647.82	2226.83
0.75	12.71	12.14	38.61	41.07	2472.50	2249.64
1.0	12.87	12.54	38.12	39.2	2271.07	2086.03

Table 3. Comparison of average vehicle speed, travel time and fuel usage for S4 with MMLC and Platooning using the Krauss and IDM car following models when varying fraction of CAVs (left most column CAVs %) equipped with MMLC and platooning in the mixed traffic network

	CO(mg/s)		CO_2(mg/s)		NO_x (mg/s)		HC(mg/s)		PM_x(mg/s)	
Col. # →	1	2	3	4	5	6	7	8	9	10
CAVs % ↓	Krauss	IDM	Krauss	IDM	Krauss	IDM	Krauss	IDM	Krauss	IDM
0	92375.84	92642.49	7067583.43	5535542.28	2779.76	2087.78	584.25	534.92	112.99	77.04
0.25	77942.21	63624.51	6545741.25	5068994.26	2535.93	1871.70	502.84	390.87	100.06	65.44
0.5	69116.71	69226.46	6159747.78	5180336.87	2352.77	1921.9	450.49	418.13	90.17	68.25
0.75	58651.23	70428.03	5751894.08	5233382.17	2147.07	1941.80	387.63	423.92	78.24	69.15
1.0	49760.34	51474.34	5283261.4	4852791.95	1920.28	1773.47	330.43	327.04	65.73	60.91

Table 4. Comparison of emissions for S4 with MMLC and Platooning using the using the Krauss and IDM car following models when varying fraction of CAVs (left most column CAVs %) equipped with MMLC and platooning in the mixed traffic network.

In Table 3, we evaluate the effectiveness and efficiency of MMLC by scaling up the network to a mixed traffic environment using strategy S4 consisting of 400 vehicles and simulated for 4750 seconds. To simulate the mixed-traffic conditions, we generated random trips for CAVs and non-CAVs. As before, when using S4 all the routes that involved the blocked link L4 are replaced by the MMLC recommended routes.

We note that only CAVs and CAVMNGRs in the mixed traffic environment are equipped with MMLC-based coordination among agents as well as platooning. Consequently, we see that as the ratio of CAVs increases to 1.0, the performance with respect to average speed increases by 2 to 6% while average travel time decreases upto 6% for the Krauss model and 9% for the IDM model. The fuel consumption monotonically decreases to about 25% reduction for the Krauss model and to about 12% reduction for the IDM model with the CAV ratio set to 1.0. Further, in Table 4, we see that the various emissions also decrease between 7% and 45% with increasing ratio of CAVs. This indicates the advantage of MMLC-enhanced CAVs in larger scale mixed traffic environments.

5. Conclusion

As CAVs become more prevalent, there is an opportunity to alleviate congestion by building smart traffic infrastructure of the future. Research on multiagent systems has the potential to contribute to this agenda. In this work, we argue that in open and dynamic environments, multiagent meta-level control is an effective way to determine when the adaptation process should be performed and how much effort should be invested in adaptation as opposed to continuing with the current action plan.

We have outlined some issues that arise in the context of an urban network – where heterogeneous actors interact – and a potential solution approach. This solution is characterized by a supervisory framework for organizational control to coordinate a vehicular network, while coping with traffic interruption events. Each acting agent consists of two levels of reasoning, one of them being a meta-level control that can reason about alternate methods for deliberation.

We have shown the advantages of our approach by means of a motivating example, in which there are CAVs that are heading to a lane that is blocked. We compared the performance of our approach, which has the ability to choreograph coordination, to situations where such coordination is not available. In a microscopic simulation for the case study presented in the paper, we were able to show that coordination led to about a 44% decrease in travel time, 7% increase in average speed, a 32% decrease in fuel consumption and about a 35% drop in emissions. Further, we have evaluated the efficiency of our approach when there is mixed traffic. Due to the fact that our approach is able to deal with rerouting (also within platoons), the advantage of having as many CAVs as possible becomes evident.

Our long term goal is to extend the investigation to other scenarios that involve coordination of further components (bikes, pedestrians, mobility-on-demand vehicles, traffic signals), which would require automated learning and self-introspective capabilities in agents while ensuring fairness and privacy within the system. We will use real data of rush hour traffic flow in New York City and simulate the behaviour of a range of actors including buses, pedestrians, delivery bicycles, CAVs and non-CAVs.

Acknowledgements

We thank Nigel Ferrer for his contribution to the MMLC framework formulation and the anonymous reviewers for their helpful comments. Support for this project was provided by a PSC-CUNY Award, jointly funded by The Professional Staff Congress and The City University of New York. Ana Bazzan is grateful to CNPq's grant 307215/2017-2.

References

Au, T.C., Shahidi, N., Stone, P., 2011. Enforcing liveness in autonomous traffic management, in: Burgard, W., Roth, D. (Eds.), Proceedings of the Twenty-Fifth Conference on Artificial Intelligence, AAAI Press. pp. 1317–1322.

Cheng, S., Raja, A., Lesser, V.R., 2013. Multiagent meta-level control for radar coordination. Web Intell. Agent Syst. 11, 81–105.

Dresner, K., Stone, P., 2004. Multiagent traffic management: A reservation-based intersection control mechanism, in: Jennings, N., Sierra, C., Sonenberg, L., Tambe, M. (Eds.), Proc. of the International Joint Conference on Autonomous Agents and Multiagent Systems, IEEE Computer Society, New York, USA. pp. 530–537.

Gordon, A., Hobbs, J., Cox, M., 2008. Anthropomorphic self-models for metareasoning agents. AAAI Workshop - Technical Report doi:10.7551/mitpress/9780262014809.003.0019.

Hasan, M.R., Bazzan, A.L.C., Friedman, E., Raja, A., 2016. A multiagent solution to overcome selfish routing in transportation networks, in: 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC), pp. 1850–1855. doi:10.1109/ITSC.2016.7795856.

Heinovski, J., Dressler, F., 2018. Platoon formation: Optimized car to platoon assignment strategies and protocols, in: 2018 IEEE Vehicular Networking Conference (VNC), pp. 1–8. doi:10.1109/VNC.2018.8628396.

Khayatian, M., Mehrabian, M., Andert, E., Dedinsky, R., Choudhary, S., Lou, Y., Shirvastava, A., 2020. A survey on intersection management of connected autonomous vehicles. ACM Trans. Cyber-Phys. Syst. 4. URL: https://doi.org/10.1145/3407903, doi:10.1145/3407903.

Klügl, F., Bazzan, A.L.C., 2021. Accelerating route choice learning with experience sharing in a commuting scenario: An agent-based approach. AI Communications 34, 105–119. doi:10.3233/AIC-201582.

Koutsoupias, E., Papadimitriou, C., 1999. Worst-case equilibria, in: Proceedings of the 16th annual conference on Theoretical aspects of computer science (STACS), Springer-Verlag, Berlin, Heidelberg. pp. 404–413.

Lopez, P.A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.P., Hilbrich, R., Lücken, L., Rummel, J., Wagner, P., Wießner, E., 2018. Microscopic traffic simulation using SUMO, in: The 21st IEEE International Conference on Intelligent Transportation Systems, IEEE. URL: https://elib.dlr.de/124092/.

Mahmassani, H.S., 2016. Autonomous vehicles and connected vehicle systems: Flow and operations considerations. Transp. Sci. 50, 1140–1162. doi:10.1287/trsc.2016.0712.

Roughgarden, T., Tardos, É., 2002. How bad is selfish routing? J. ACM 49, 236–259.

Sala, M., Soriguera, F., 2020. Macroscopic modeling of connected autonomous vehicle platoons under mixed traffic conditions. Transportation Research Procedia 47, 163 – 170. doi:https://doi.org/10.1016/j.trpro.2020.03.089. 22nd EURO Working Group on Transportation Meeting, EWGT 2019, 18th - 20th September 2019, Barcelona, Spain.

Treiber, M., 2013. Traffic Flow Dynamics. doi:10.1007/978-3-642-32460-4.

Vasirani, M., Ossowski, S., 2009. A market-based approach to reservation-based urban road traffic management, in: Decker, K., Sichman, J., Sierra, C., Castelfranchi, C. (Eds.), Proc. of the 8th Int. J. Conf. on Autonomous Agents and Multiagent Systems (AAMAS), IFAAMAS, Budapest. pp. 617–624.

Wang, C., Gong, S., Zhou, A., Li, T., Peeta, S., 2019. Cooperative adaptive cruise control for connected autonomous vehicles by factoring communication-related constraints. Transportation Research Procedia 38, 242–262. URL: https://www.sciencedirect.com/science/article/pii/S2352146519300237, doi:https://doi.org/10.1016/j.trpro.2019.05.014. journal of Transportation and Traffic Theory.

Yamane, S., Ishida, T., 2006. Metal-level control architecture for massively multiagent simulations, in: Proceedings of the 2006 Winter Simulation Conference, pp. 889–896. doi:10.1109/WSC.2006.323172.