Multimedia Tools and Applications (2021) 80:15921–15936
https://doi.org/10.1007/s11042-020-10222-y

About auction strategies for intersection management when human-driven and autonomous vehicles coexist

Giacomo Cabri[1] · Luca Gherardini[1] · Manuela Montangero[1] · Filippo Muzzini[1]

Received: 21 December 2019 / Revised: 25 March 2020 / Accepted: 9 December 2020 /
Published online: 6 February 2021

Abstract

Autonomous vehicles are appearing in our streets, and will soon populate our transportation infrastructures, which must be equipped with appropriate sensors and actuators in order to manage vehicles in a fruitful way. Besides the infrastructures, appropriate algorithms must be defined in order to coordinate the vehicles and to enable them to exploit the resources in a fair yet effective way. In the immediate future, autonomous vehicles must coexist human-driven vehicles, and this transitory scenario poses several challenges in coordinating both kinds to exploit street resources. One of these resources, whose management is quite challenging, is represented by intersections: vehicles come and aim at passing the intersection, often as soon as possible, but they must compete with other vehicles having the same aim. A possible approach that has been used in literature to this problem uses auction based mechanisms. In this paper, we place ourselves in the above-mentioned transitory scenario in which both human-driven and autonomous vehicles will compete to cross intersections, and we investigate the effectiveness of auction-based mechanism to coordinate vehicles at intersections. We devise some simple auction policies, and assume vehicle coordination strategies that are suitable also for human drivers. Our results lead us to believe that, under these assumptions, simple auction mechanisms do not introduce advantages for what concern traveling times as they do in the case of exclusively autonomous vehicles.

Keywords Autonomous vehicles · Coordination · Auction

✉ Manuela Montangero
manuela.montangero@unimore.it

Giacomo Cabri
giacomo.cabri@unimore.it

Luca Gherardini
218121@studenti.unimore.it

Filippo Muzzini
filippo.muzzini@unimore.it

[1] Dipartimento di Scienze Fisiche, Informatiche e Matematiche, Università di Modena e Reggio Emilia, Via Campi 213/B, 41125 Modena, Italy

1 Introduction

With the rapid growth of the IoT [6] and new opportunities of *making almost everything smart*, smart mobility [18] will improve either safety on our streets and our experience of driving in smart cities: our streets will be soon populated by autonomous vehicles, which will take autonomous decisions about the route to follow, not only considering the fastest path between the origin and the destination, but also considering the safest driving conditions, the cost of the route, the needs of the driver, the current traffic, the proximity of cultural heritage locations and so on [5].

However, smart mobility must be supported [7], not only by an appropriate infrastructure that connects vehicles and the environment, but also by appropriate algorithms that manage and coordinate the vehicles' movements. In particular, *intersections* represent a challenge [3]: traditional traffic lights and yield rules have been introduced to coordinate vehicles at intersections keeping in mind characteristics of both human drivers (such as unpredictable behaviors and reaction times) and intersections themselves (leading to decisions such as giving yield to the larger and highly trafficked lane).

The scenario changes completely when vehicles are autonomous and have the possibility to communicate among themselves and acquire information about the surrounding environment. Any number of alternative solutions might be devised with the goal of achieving a better traffic coordination, by overcoming the rigidity of traditional traffic lights in assigning resources (i.e., the intersection), and the starvation phenomenon in which vehicles driving on roads without yield might incur in. Some approaches were proposed to address collaborative management of intersections [12, 22]. Often, auctions are exploited to dynamically manage resources (e.g., see [11]), because they enable to define the price of a good on the base of customer interest [2]. Auctions were largely exploited to manage negotiations between autonomous entities modeled as agents [10]. However, solutions designed for a scenario in which all vehicles are autonomous might not be suitable for the unavoidable (who knows how long) transition period in which human-driven and autonomous vehicles will coexist. In particular, we can not assume human drivers to be reliable, willing or able to follow given instructions. For example, at intersections, we can assume that human drivers might execute simple instructions as "stop and wait for your turn" or "go, it's your turn". On the contrary, we can not assume that they will execute more detailed sequences of instructions such as, for example, "slow down to 35 km/h and proceed for 10 meters, slow down to 27 km/h and proceed for 5 meters, accelerate to 32 km/h and turn 35 degrees on the right...". Analogously, we can not ask human drivers to execute very precise instructions such as, to pass close to other vehicles with a few cm of precision. Finally, it is unthinkable to make human drivers substitute sensors and cameras to collect data of the surrounding environment in order to implement coordination policies, such as adjusting speed according to proximity with other vehicles. Therefore, in this paper we will propose coordination policies and auction mechanisms that take into consideration limitations induced by humans driving vehicles.

The present work is part of a longer-term project in which we aim to devise coordination policies for coexisting human-driven and autonomous vehicles [8, 9]. In this paper we present experiments related to the case in which coordination is achieved by means of auctions, taking into consideration design limitations due to the presence of human-driven vehicles. At this aim, we devise some simple auction strategies and we analyze latencies experienced by vehicles at intersections. Our results lead us to the unexpected conclusion that simple auction strategies are not so effective in our scenario, especially in the most critical situation of heavy traffic conditions.

The paper is organized as follows. In Section 2 we present and discuss some related work, then, in Section 3 we present the approaches we propose to manage vehicles at crossings. Experimental results are shown and discussed in Section 4. Finally, we conclude the paper proposing some future work (Section 5).

2 Related work

There are different approaches related to the coordination of autonomous vehicles. We will first present general approaches in Section 2.1. Then, in Section 2.2 we will focus on approaches for coordinating vehicles at crossings: there might be different kinds of solutions, from the centralized ones to the most distributed ones, where a global coordination emerges from the behavior of the single vehicle, without neither a centralized manager nor interaction among vehicles. In this section, we will focus on approaches based on auctions, which is the mechanism we exploit in our proposal.

2.1 General coordination of autonomous vehicles

In literature, it is possible to find reports of many research efforts related to the design of the future smart cities (see [17] for a survey). In such scenarios, autonomous vehicles will have to interact with their surrounding environment: this will enable the software residing in each vehicle to take decisions in presence of other vehicles. Such decisions will also be influenced by a newly designed intelligent infrastructure for smart cities. Even if such topics are not new in theoretical research, it is just recently that newly released compute accelerators have reached SWaP (Size, Weight and Power) features to allow for experimenting on medium to large scale fleet of intelligent vehicles in specially equipped city areas [14, 16].

Pinciroli et al. [20] take into consideration the programming of autonomous robots. They highlight the difference between smart devices and autonomous vehicles in the scenario of navigation; the first ones have small capability of interacting with the physical world, while this capability is exhibited by the robots in general and by autonomous vehicles in particular. To this purpose, the latter can exploit sensors and actuators, which allow also to act on the surrounding environment. Their proposal relies on a *swarm language construct* that allows to categorize robots in swarms and to assign jobs to the swarms. The aim of their approach is to provide for *re-usability* and *predictability* of the coded behavior, which turn out to be very important issues in the field of autonomous driving.

Murthy et al. [19] propose a simulated environment designed to have cars traveling in a highway, self-organize themselves in platoons with the final goal to reduce fuel consumption, by drafting off one another. This approach is in the same direction of the previous one, in order to achieve *re-usability* and *predictability* in a specific application scenario.

Our work is framed in the context of the CLASS Horizon 2020 project (https://class-project.eu/), whose goal is to arrange a smart area of one square Kilometer-wide in the city of Modena (Italy). Thanks to a large sensor infrastructure, we are collecting and processing in real-time the resulting vast amount of data. Such data will be used to communicate to the connected vehicles. Such vehicles are equipped with heterogeneous sensors/actuators and V2X connectivity so to enhance both driving experience and city overall safety. This is achieved by deploying advanced urban mobility applications based on a combination of data-in-motion and data-at-rest analytics to efficiently coordinate cars and city computing resources. The part of the city involved is depicted in Fig. 1.

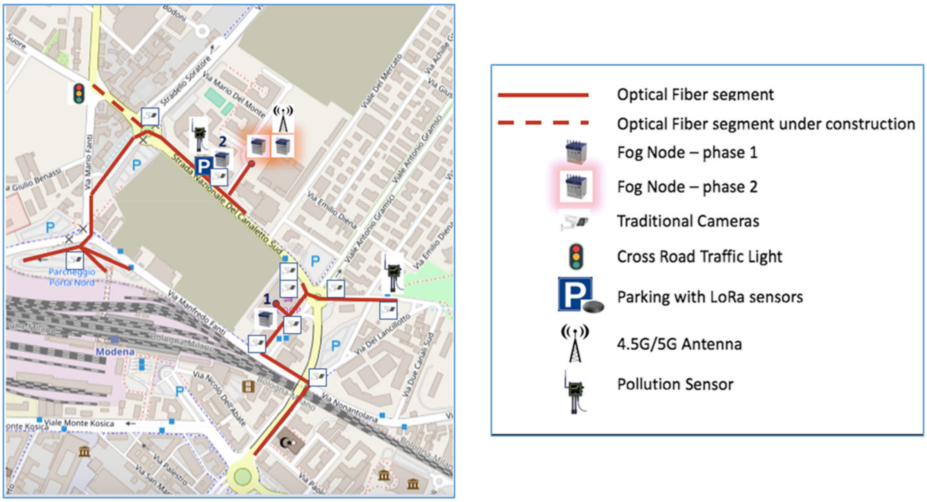

Fig. 1 Modena Automotive Smart Area. Numbers 1 and 2 indicate the positions of Fog Nodes that collect data from edge sensors. LoRa (Long Range) is a patented wireless data communication technology typically adopted in natural resource management, renewable energy and other long-range IoT devices

2.2 Intersection management by auctions

In one of the milestone papers exploiting auctions in management crossings [11], authors address the value of time of each vehicle, representing it by means of a wallet system for automatic bidding based on trip characteristics, driver-specified budget, and remaining distance to the destination. Moreover, they address also optimization of the overall traffic. The main differences with our work lies in the scenario we set ourselves in by considering coexisting human-driven and autonomous vehicles. This implies that we can not implement optimization traffic strategies that require complex vehicle coordination.

A more complex approach is adopted by Schepperle and Böhm [23]. They propose a two-step auction mechanism to manage vehicles at intersections, based on a second-price sealed-bid auction [25]: in the former step only vehicles that can pass the intersections are involved, while in the latter step also some vehicles in the second place of the lane can bid, depending on the result of the first step for the preceding vehicle. In the first step, following the second-price auction rule, the system allocates the next time slot to the vehicle with the highest bid, while the vehicle needed only to pay the second highest bid. In the second step, only vehicles with their previous vehicles having achieved time slots can bid and only one vehicle from each lane was allowed to bid. Even if that approach was quite effective, we decided to define a *simpler* approach, in order to consider the total number of vehicles in line to enhance the bid of the vehicle in the front.

Finally, Vasirani and Ossowski [24] propose a different approach to manage urban crossings, based on a reservation-based intersection control model, proposed in [15] and on market-inspired rules. The authors consider two scenarios: one with a single intersection and one with a network of intersections. The first is exploited to analyze the performance of a policy based on combinatorial auctions to allocate slot reservations. The second is exploited to analyze the impact that a traffic assignment strategy inspired by competitive markets has on the drivers' route choices. They combine the two to propose an adaptive management mechanism based on the auction-based traffic control policy and the competitive traffic assignment strategy. Differently from that approach, we do not assume that all vehicles are autonomous, considering also human-driven vehicles for which we can not make any assumption concerning the knowledge of the vehicles routes.

More in general, auctions have been exploited also to manage resources different from intersections, such as parking slots and fleets. Indeed, parking is a highly competitive resource allocation coordination problem, when considering private and individual vehicles. Among different proposed solutions (see [21] for a survey), auctions are used to implement a negotiation approach in which each vehicle actively participate to the coordination process advancing its proposals [1, 13]. As for fleets of special vehicles, like emergency vehicles or taxis, they usually have a limited number of units. Thus, to handle cases in which the number of users exceeds the number of units, we need a policy to decide which users to serve first, and auctions have been exploited for this purpose [4].

We can highlight two main differences between this body of work and our contribution. First, we consider both *autonomous* vehicles and *human-driven* vehicles. Handling such a simultaneous presence is of paramount importance for coping with the transition from only human-driven vehicles to only autonomous vehicles. Indeed, it is considered in one of the foundation papers on autonomous intersection management [15]. However, such copresence imposes to take into consideration several constrains related to the fact that human driven vehicles are not predictable and not completely trustful.

Our contribution shares with the last paper a second important difference from the above proposals on auction-based management: it lets *all* vehicles waiting in a lane to participate in the auctions, and not only vehicles at the head of a lane.

3 Our proposal

In our scenario we assume that each intersection has its own management system that is run on a physical device that is placed at intersection site. The vehicles crossing priorities are decided by the intersection management system by means of an auction: vehicles at the front of each lane make their own bids, whose amount is defined according to their budget. Moreover, the front lane vehicle's bid might be enhanced by taking into consideration the number of vehicles in the lane, or either by allowing vehicles in the back of the lane to contribute to the bid.

We further assume that autonomous vehicles and human-driven vehicles coexist and that all vehicles are equipped to participate to the crossing management mechanism. In particular, human-driven vehicles do not require humans to directly intervene in the auction, but vehicles will be equipped and autonomously able to participate to the auction mechanism. We observe that this assumption is not so unreasonable or difficult to achieve. Indeed, mobile phone Apps or systems as Infotainment might be exploited: they are activated at departure, they will manage all communications with the management system related to auctions, and will be used to instruct drivers. Consequently, from the auction point of view there is no difference between human-driven and autonomous vehicles. However, we will not be able to implement any coordination mechanisms that assumes knowledge of vehicles trajectories or in general that exploits forecast of vehicles behaviours, simply because human-driven vehicle behaviours are unpredictable and can not even be trusted to follow a suggested or predetermined route. In particular, we can not implement any policy to make clearing of intersection faster (or smarter) by allowing two or more vehicles to access the intersection at the same time if their trajectories are not conflicting (resolution adopted for example in [11]), either because human driven vehicle trajectories are not known or because we can not trust human driver to actually follow established ones.

Therefore, in the following we will assume that only one vehicle at the time is allowed to access the intersection, and that the only instructions given to human drivers, when they get at the beginning of a lane, are either to *stop* or *drive through* the intersection.

In this paper we aim to study very simple auction mechanisms with a twofold goal in mind, naturally arising by looking at the benefits of auction mechanisms for autonomous vehicles. The first is to show if auction mechanisms might contribute to reduce vehicle latencies at intersections, with respect to standard traffic yield rules. The second is to understand if and how to achieve differentiated latencies; i.e., higher bids resulting in smaller latencies. Indeed, as lanes are cleared following a FIFO policy (to approach an intersection vehicles have to wait that all other vehicles preceding them in the lane have crossed it), it is not trivial to conclude straight away that just placing high bids vehicles are able to reduce their latencies.

In the following of this section we introduce the mechanisms that we will study in our experiments.

3.1 Bidding strategies

We assume that bids are set using virtual coins and that vehicles get a certain budget (in terms of such virtual coins) as they join the system.

We are not concerned about how vehicles get their initial budgets, as this is not the focus of the paper. In real scenarios, virtual coins will be implemented by means of different policies according to rules defined by local administrations, taking into account different social, economical and/or ethical aspects. For example, virtual coins might be bought with real money, or might reflect the drivers (good or bad) behaviour while driving, e.g., respect of driving rules, adoption of car pool, reduction of car usage, choice of low emission or green engines, and so on. Therefore we think that, rather that technical, deciding how to implement virtual coins and budgets, is an ethical and political issue concerning suitable mobility regulation and avoiding mobility becoming a privilege.

We design two different bidding strategies:

- **Randomized (Rand):** vehicles place a random bid that falls inside the range of their budgets. If a vehicle runs out of budget, the system will let it make a minimal bid anyway.
- **Route dependent (Prop):** vehicles divide their budgets by the total number of intersections they have to go through during their trip, and they use the result to place bids at each intersection. We observe that, if routes are fixed and do not change during the trip (as for example due a reroute for heavy traffic), vehicles should not run out of budget. To make human driven vehicles adopt this strategy, human drivers will be required to fill in their final destination in their navigation system at departure. The auction mechanism will then compute the shortest route to go from the starting point to destination, it will consider this route the expected route, and will compute the bid amount according to the number of intersection in the expected route. If the driver follows a different route, a new expected route is computed and the bid amount adjusted consequently.

In our experiments we will assume that virtual coins are indivisible, thus bids will be integer numbers. However, if in a real scenarios virtual coins are divisible, bids might be drawn with a precision that reflects the number of digits of sub-multiples of virtual coins.

We also devise a *budget recharge* mechanism in which the amount of the bid placed by the winner of an auction is equally redistributed among the other vehicles that participated to (and lost) the auction. The mechanism is especially thought for the Rand bidding strategy, to avoid vehicles to run out of budget early in their trip, however we will test it also for the Prop bidding strategy. We are aware that it might seem unnatural or unfair to redistribute some vehicle's budget to others (especially if, for example, budgets is bought with real money), however if there is a common benefit in doing so, then it is a mechanism worth adopting.

3.2 Auction resolutions

We devise two different approaches for auction resolution:

- **Cooperative approach (COOP):** all the vehicles at the front of the lanes at the intersection make their bid, and all bidding vehicles will go through the intersection according to the bid order (highest first).
- **Competitive approach (COMP):** all the vehicles at the front of the lanes at the intersection make their bid, but only the vehicle that wins the auction gets to pass the intersection. Vehicles that lost the auction will have to attend and win a successive auction before being able to go through the intersection. Note that this may imply *starvation*.

Moreover, we consider two different methods for the bid payment:

- **All-Pay (AP):** All bidding vehicles are charged for their bid.
- **Only-Winner-Pays (OWP):** only the vehicle that wins the auction will be charged for its bid, while vehicles that lost the auction are not charged.

In a preliminary study presented in [9], we investigated the effectiveness of the four strategies resulting from all four possible combinations of the above auction resolution approaches and payment methods, when using the simple randomized bidding strategies. We concluded that two out of four combinations were more interesting than the others, namely strategies COOP-AP and COMP-OWP, those more natural to interpret. In the following of the paper we will focus only on these two strategies. These are complementary: in the former all vehicles have to pay their bids, but in exchange they get to go through the intersection after the auction is over, they just do not know in advance in which order. This round-robin like approach guarantees that all lanes proceed at each auction round. In the latter, on the contrary, lanes moves at different speed depending on which lane the winner of the auction is in. We measured average times to clear intersections (i.e., elapsed time from the moment vehicles reach the front of the lane and the moment they exit the intersection) and average waiting time in lane, under different traffic conditions. Results showed that average time to clear intersection is very small for COOP-AP, and slightly more than doubles for COMP-OWP; variances of average values are small for both cases. For what concerns average waiting time in line (and corresponding variances), there is no drastic difference between COOP-AP and COMP-OWP under the same traffic conditions, the latter being slightly larger than the former.

Finally, only for the competitive approach, we devise two mechanisms that allow other vehicles that are not at the beginning of a lane to be involved in the auctions: *enhancement* and *sponsorship*, described as follows:

Enhancement The *enhancement* mechanism is meant to balance the waiting times of lanes with long lines with respect to lanes with a smaller number of vehicles in line. To this aim, we "adjust" the front lane bidding according to the number of cars waiting in the same lane. In particular, for each lane ℓ with n_ℓ cars in line and bid b_ℓ made by the car at the front of the lane, we formally define the enhancement in the following way:

$$en(\ell) = b_\ell(\ln n_\ell + 1),$$

and the value $en(\ell)$ is the actual bid for the front vehicle of lane ℓ. Observe that $en(\ell) = b_\ell$ when the only vehicle in the lane is the front one, while the logarithmic multiplicative factor gets the more significant the larger the number of vehicles waiting in the lane. In this mechanism, vehicles involvement in the auction is implicit, in the sense that they do not place a bid, but their presence in a lane influences the bid of the vehicle in the front of the lane.

Sponsorship The *sponsorship* mechanism allows vehicles behind in lanes to place bids to be added to the bid of the vehicle at the front of the lane. The sponsorship bids will be charged only if the sponsored vehicle wins the auction. In this mechanism, vehicles involvement in the auction is explicit, and is intended to allow vehicles to speed up their lanes at the expenses of the others.

We will not exploit enhancement and sponsorship in the cooperative approach because, with such an approach, for each lane, one vehicle gets always to pass the intersections at each auction, even if it looses the auction. On the contrary, in the competitive approach lanes might incur in starvation and these mechanisms are thought to avoid this situation to occur.

4 Experiments

In this section we give details of our experiments and we discuss their results.

4.1 Implementation and experimental setup

We conducted our experiments using two different simulators, namely SUMO[1] (as in our preliminary work [9]) and MATSim urban simulator,[2] and we comfortably got comparable results. Here we report results obtained using the latter simulator.

We used a map that is a Manhattan grid with 8×8 intersections, in which each link has exactly one lane for each direction. Vehicles follow random routes: starting and ending route location have been selected randomly on the map and the route is computed as the shortest path from starting to ending location. Departure times follow two Gaussian's distributions, one with its peak at 9AM and the other at 6PM, to simulate daily traffic. Each vehicle has an initial budget set to a default value. We run experiments with a different number of vehicles to simulate different traffic conditions.

[1] Simulation for Urban MObility, https://www.dlr.de/ts/en/desktopdefault.aspx/tabid-9883/16931_read-41000/

[2] https://www.matsim.org/

Each experiment is characterized by its own scenario, defined by selecting one approach for auction resolution (COOP-AP or COMP-OWP) and one bidding strategy (Rand or Rd), with and without budget redistribution. Moreover, for the COMP-OWP approach, we separately analyzes cases in which enhancement and sponsorship are enabled or disabled (respectively Enh and Spon). We also run experiments in which only standard yield rules are applied (referred to as Priority in Figures). Waiting times are measured in seconds.

For each experiment, we report the following measures:

– Statistics on vehicles *average waiting times at intersection* before they are given the right to pass the intersection: from the moment vehicles reach the front of the lane, to the moment they win an auction, in the COMP-OWP case, or to the moment it is their turn to pass trough after the auction is over, for the COOP-AP case.
– Statistics on *average times vehicles spend waiting in lanes*: we compute the average time each vehicle spends waiting in lane, from the moment it queues up, to the moment it has cleared the intersection.

If a vehicle destination lane is filled up, we make the vehicle wait at the front of its present lane before crossing the intersection, and no other vehicle is allowed to go trough (equivalently we could have made the vehicle stop in the middle of the intersection, blocking the passage to other vehicles). Indeed, in our scenario this is the safer solution. One might think, for example, of alternative solutions in which other vehicles are allowed to pass the intersection their destination lane is not filled up, even if they did not win the auction. This might work when vehicles are all autonomous, but in our scenario there is no way for the intersection management system to know in advance the trajectories of human-driven vehicles. Moreover, making on-the-fly changes on the order in which vehicles are allowed to go trough intersection might require to give several contrasting indications to human drivers, increasing the chances that they make mistakes and cause incidents.

4.2 Experimental results

In this section we report the results of the experiments conducted as explained in the previous section. In Section 4.2.1 we present results of the comparisons among the different scenarios and latencies experienced by adopting standard yield rules. In Section 4.2.2 we present results related to achieving differentiated latencies.

4.2.1 About reducing latencies

We conducted simulations with 5K vehicles running on the map simulating heavy daily traffic conditions. Figure 2 reports the average of the waiting times for vehicles at the front of a lane, i.e., average time before being given the right to cross the intersection. We show the two cases in which we adopt or do not adopt the budget recharge mechanism. We start observing that with any combination of our strategies, on average, vehicles gain their turn faster than when adopting standard yield rules (Priority in the figures).

As expected (and coherently with our preliminary experiments in [9]), among all auction policies, COOP-AP is the one that guarantees smaller waiting times, under all different alternatives. The reason why is very simple: vehicles participate only to one auction and they will never have to wait more than one vehicle per each lane of the intersection. Indeed, we observe a very small variance (black vertical line in the plots) for all COOP-AP cases. When considering the COMP-OWP case, the difference between redistributing or not the budget is evident only in absolute times (smaller in the former case), but the relative behaviour of

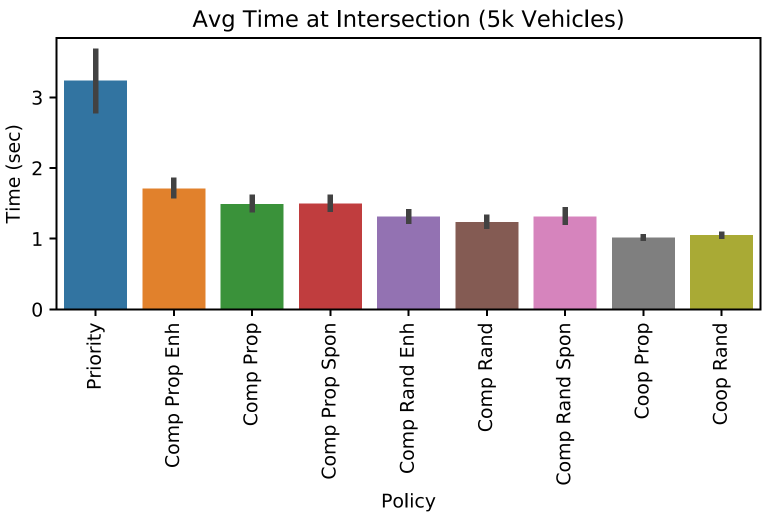

(a)

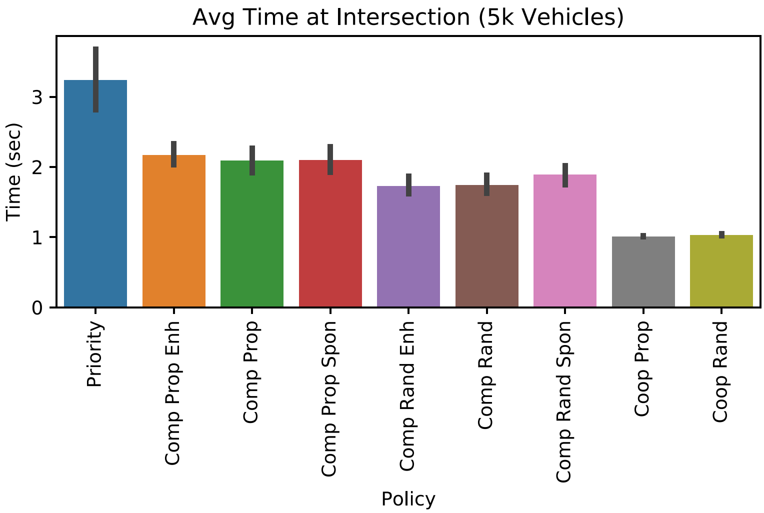

(b)

Fig. 2 Average waiting times at intersection before given the right to pass the intersection, with 5k vehicles running on the map. Vertical black lines represent standard deviation. **a** Budget recharge mechanism is adopted. **b** Budget recharge mechanism is not adopted

different policies is practically the same. Times achieved with the Rand bidding strategy are slightly smaller than when using the Prop strategy, as in the former the chance of losing several auction in a row is lower. This happens because bids of the same vehicle ranges in the whole budget interval and chances are that a small bid (causing the vehicles to loose one auction) is followed shortly after by a much larger one (causing the vehicle to win the auction). On the contrary, with the Prop strategy, vehicle bids are always very similar one to the other. Finally, when adopting *enhancement* and *sponsorship* we can observe a slight increase of average times with respect to non adopting them. This happens because the number of vehicles that experience lower latencies in the former case are in a smaller number than those who experience a higher one.

We also measured total waiting times on lanes, i.e., from the time vehicles queue-up, to the time they clear the crossing by entering the next lane, when budget is redistributed and not redistributed. Results for the two cases show the same behaviour. Unexpectedly, any benefit exhibited when evaluating only auction times at intersection, is completely lost when one considers the larger picture: any strategy incurs in practically the same average (and standard deviation) and maximum times for vehicles to go from one lane to the successive. There are only insignificant differences that are not even evident when plotted, as too small compared to the used scale. Therefore, we do not provide plots, as they are meaningless. Our understanding of this negative result is that lane latencies are more conditioned by general traffic conditions than by how fast one vehicle gets the right to go trough intersections by participating to auctions.

Hence, in the subsequent experiments we varied the number of vehicles running on the map. Figure 3 shows plots related to average traffic conditions (2500 vehicles), while Fig. 4 show plots related to light traffic conditions (200 vehicle). We can see that the general picture does not change, except for absolute values (smaller for lighter traffic).

With light traffic we have a slightly larger standard deviation, meaning that there is a larger gap between the best and the worst waiting time in the different traffic configurations that might occur at intersections (i.e., different number of lanes with different number

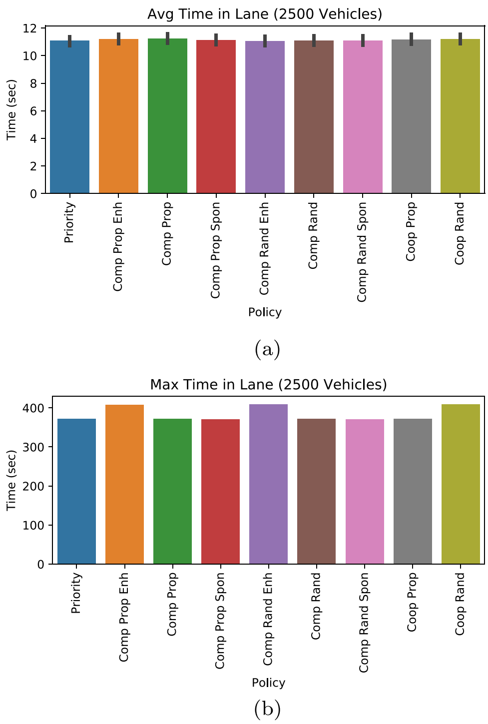

Fig. 3 (**a**) Average waiting times spent waiting in lane, with 2500 vehicles. (**b**) Maximum time spent waiting in lane, with 2500 vehicles. In both cases budget is not redistributed

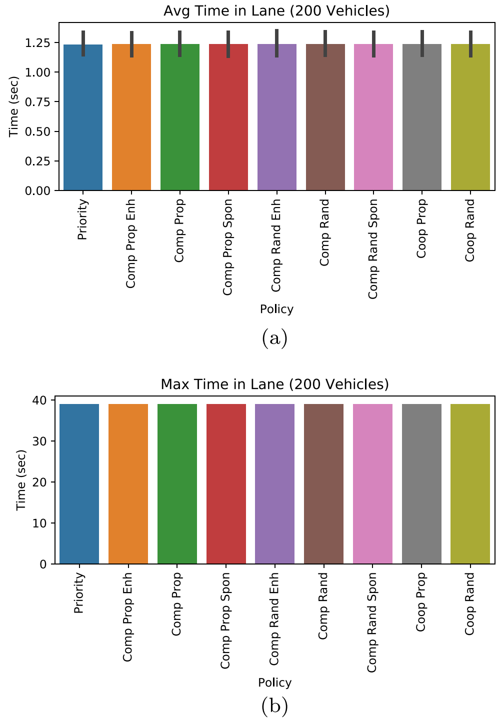

Fig. 4 (**a**) Average waiting times spent waiting in lane, with 200 vehicles. (**b**) Maximum time spent waiting in lane, with 200 vehicles. In both cases budget is not redistributed

of vehicles). However, on average such differences resolves in similar times, as there are very few vehicles (and evidently any policy works equally well). Analogously, all policy incur in comparable worst cases, i.e., maximum waiting times. This result can be explained by observing that, with light traffic vehicles, a small number of different configurations happens at intersections, and evidently the same worst case happens at least once in each scenario.

As for the average traffic condition setting, we observe only some differences in maximum waiting times, that are not difficult to explain remembering that such values might be experienced by one (or very few) vehicles. Indeed, COOP-AP-Rand either once experienced one longer lane than COOP-AP-Prop, or in the lane the front vehicles placed a large number of small bids. On the other side, for the COMP-OWP with *enhancement* larger maximums records the fact that one vehicle had to wait longer, reasonably because some other lane was longer and advantaged by the mechanism. When considering average waiting times, we observe practically the same average independently from the specific adopted scenario. Our understanding of this result is that, when traffic is heavy, waiting time is dominated, on average, by the time needed to clear the full lane ahead the vehicle (and as in our experiments the map is a grid, all lanes have the same length).

We wish to point out that we got analogous results also from the other set of experiments, conducted using a different simulator, implemented independently, and using a smaller 3×3 Manhattan style map. In particular, for total waiting times in lanes, we observe some differences (coherent with the plot in Fig. 3b) only for maximal waiting times, in the case of average traffic conditions (one hundred vehicles).

We conclude by observing that any of our strategies do not give significant improvements for what concerns latencies, when compared to adopting standard yield rules. Results thus suggest that simple auction mechanisms might not be a suitable coordination policy to reduce vehicles latencies in the transition period toward exclusively autonomous vehicles, because traffic flow and lane overload affect waiting times more than the advantages introduced by auctions.

The only advantage that we see concerns traffic coordination and the possibility to reduce the number of accidents among vehicles. Indeed, auction mechanisms are able to clearly indicate which vehicle has the right to get through the intersection eliminating misunderstandings and miscalculations that commonly lead to accidents among human-driven vehicles under traditional yield rules. However, auctions are just one possible solution to reach this last goal.

4.2.2 About differentiating latencies

In this section we take into consideration the specific case in which one would like to make one single vehicle be able to definitively speed up its lane, as when the vehicle is an emergency one, e.g., an ambulance or a fire engine. What happens nowadays, on our urban streets populated mainly by human-driven vehicles, is that emergency vehicles announce their urgency using colored flashing lights and acoustic signals. The other vehicles on the street try to pull over to make it pass. We wish to understand if it possible to exploit the auction mechanism as is, in order to allow emergency vehicles to speed up their lane, without the need to implement specific ad-hoc rules for such cases.

Therefore, we set up experiments in which we have only one vehicle with a possibly infinite budget, allowing it to place extremely large bids that are non comparable to any other vehicle bids. We measure the difference of waiting times for the vehicle when endowed with infinite budget and when given standard budget, always traveling at the same time of the day. We make different measurements by varying traffic conditions and emergency vehicle trajectory. In particular, one trajectory is crocked and goes through many intersections, the other is a straight route from one side of the map to the opposite one.

Our finding is that we do not reach the desired result under any combination of strategies. In the following we analyze the reasons why we got such results.

Cooperative approach Under the COOP-AP assumption we do not see any significant advantage in giving the emergency vehicle such a high budget; however this result is not surprising. This policy is extremely fair, and there is no way for a single vehicle with high budget to be able to speed up its own lane. If lucky, its lane will always be the first to clear the intersection. Nevertheless, before the next auction, all vehicles must wait that all vehicles at the lane fronts clear the intersection.

To be more precise, let L be the number of lanes at the intersection, Δ the average time needed by a vehicle to clear the intersection and p the position of the emergency vehicle (or any other vehicle) in its lane. Consider the worst case in which all lanes are crowded, at least as much as the lane the emergency vehicle is in. Then, in order to reach the front of its lane, in the best case in which its lane always wins the auctions, the emergency vehicle (or any other vehicle) has to wait at least $\Delta \cdot L \cdot p$ time units, independently of its current budget.

Competitive approach Under the COMP-OWP we expected to see evidence of benefits of giving a high budget to the emergency vehicle. However, we never observed significant improvements. The first observation we have is that, when traffic is light, waiting times are low in general because many links are free and, if not free, lines are always very short. As also the emergency vehicle has to wait for vehicles ahead in the lane to pass the intersection before it, having a high budget is not incisive in getting significant lower waiting times. Therefore, we concentrated on heavy traffic conditions, which is the real situation in which one would like emergency vehicles to be able to move faster than the others.

The Rand bidding strategy resulted not interesting because, whenever the random chosen bid is small, the emergency vehicle loses any benefit in having a potential infinite budget. It is true that we expect the bid to be large enough in a small number of tries, but in the meantime the emergency vehicle loses the ability to clear its lane before others. On the contrary, the Prop bidding strategy assures that the emergency vehicle bids are always incomparable larger than the others.

Enhancement is not helping either, because it takes into consideration only the number of vehicles in the lanes and not their budget. Indeed, whenever lanes are crowded, the emergency vehicle counts as one exactly as any other vehicle in the lanes, and the winner is decided by means of the bid of the vehicle at the front of the lane. The only way in which the emergency vehicle bids are taken into consideration even if it is back in the lane is to adopt the *sponsorship* mechanisms.

Applying *budget recharge* we experienced even larger latencies for the emergency vehicle endowed with infinite budget. The simple explanation is that whenever the emergency vehicle wins the auction, all front lane vehicles get recharged with an infinite budget. From that moment on, they start acting as emergency vehicles themselves, interfering with successive auctions involving the original emergency vehicle. This negative result is particularly evident in our experiments with the crocked emergency vehicle trajectory. Indeed, we forced it to cross many intersections and, accomplice the Manhattan stile map, it happens that two high budget vehicle compete more than once at different intersections.

Finally, neither the most promising COMP-OWP approach, with Prop bidding strategy, no recharge mechanism, and *sponsorship* showed evidence of being effective. The reason is not that the emergency vehicle is not able to make its lane win all auctions, but that it does not succeed in speeding up the lane. The problems is that, with high traffic, even if the front lane vehicle wins the auction, it can not clear the intersection if its destination lane is full, and the emergency vehicle can not influence the auctions occurring at other intersections.

We conclude by observing that none of the proposed combination of bidding/auction strategies simply adapts to minimize latencies of vehicles with potential infinite budget under any traffic condition. However, the really interesting results is the reason why even the "COMP-OWP approach combined with Prop bidding strategy, no recharge mechanism, and *sponsorship*" fails. If an emergency vehicle is not allowed (or has no possibility) to overtake other vehicles in its line, heavy traffic causes the vehicle to get stuck even if it is given infinite budget. Our intuition is that this situation is independent of the specific auction policy adopted at the intersection, but reasonably depends on the fact that each intersection acts independently. Addressing this intuition will be the focus of future work.

5 Conclusions

In this paper we presented a study on several simple intersection management systems based on auctions, comparing them among themselves and against the situation in which standard yield rules are applied to coordinate vehicles at intersections. With respect to other proposals on vehicle coordination, we place ourselves in the different scenario in which human-driven and autonomous vehicles coexist. We observe that this assumption highly influences the design of coordination policies because human drivers preclude the adoption of many solutions that can be exploited when vehicles are exclusively autonomous.

We investigated the effectiveness of auction based mechanisms, as these have been shown to bring interesting benefits for autonomous vehicles: latencies reduction and the possibility to differentiate latencies among vehicles according to their needs. Our main result is a negative result, as in our experiments and our scenario, we did not experience any of these two advantages in using auctions. Indeed, limitations design introduced by human drivers frustrates the auction system benefits.

With regard to future work, we will deeply investigate motivations behind our results to better understand how much of these depend solely on the specific scenario that we considered, or on the fact that our strategies are too simple, or even on the fact that vehicles are necessarily serialized when in lanes. Thereafter, we will study alternative coordination policies for the scenario considered in this paper, namely when human-driven and autonomous vehicles coexist.

Acknowledgments This work was supported by the EU H2020 program under the CLASS project, grant No. 780622.

References

1. Alves BR, Alves GV, Borges AP, Leitão P (2019) Experimentation of negotiation protocols for consensus problems in smart parking systems. In: International conference on industrial applications of holonic and multi-agent systems. Springer, pp 189–202
2. Bergemann D, Said M (2010) Dynamic auctions. Wiley Encyclopedia of Operations Research and Management Science
3. Bertogna M, Burgio P, Cabri G, Capodieci N (2017) Adaptive coordination in autonomous driving: motivations and perspectives. In: IEEE int. conf. on enabling technologies: infrastructure for collaborative enterprises (WETICE). IEEE, pp 15–17
4. Billhardt H, Fernández A, Lujak M, Ossowski S, Julián V, De Paz JF, Hernández JZ (2016) Towards smart open dynamic fleets. In: Rovatsos M, Vouros G, Julian V (eds) Multi-agent systems and agreement technologies. Springer International Publishing, Cham, pp 410–424
5. Bujari A, Ciman M, Gaggi O, Palazzi CE (2017) Using gamification to discover cultural heritage locations from geo-tagged photos. Pers Ubiquit Comput 21(2):235–252
6. Bujari A, Furini M, Mandreoli F, Martoglia R, Montangero M, Ronzani D (2018) Standards, security and business models: key challenges for the iot scenario. MONET 23(1):147–154
7. Bujari A, Gaggi O, Palazzi CE, Ronzani D (2018) Would current ad-hoc routing protocols be adequate for the internet of vehicles? A comparative study. IEEE Internet of Things Journal 5(5):3683–3691
8. Cabri G, Muzzini F, Montangero M, Valente P (2020) Managing human-driven and autonomous vehicles at smart intersections. In: Proc. international conference on human-machine systems, to appear
9. Cabri G, Gherardini L, Montangero M (2019) Auction-based crossings management. In: Proceedings of the 5th EAI international conference on smart objects and technologies for social good, GoodTechs '19. ACM, New York, pp 183–188
10. Capodieci N, Andrea Pagani G, Cabri G, Aiello M (2011) Smart meter aware domestic energy trading agents. In: Proc. of the 2011 workshop on e-energy market challenge, IEEMC '11. ACM, New York, pp 1–10
11. Carlino D, Boyles SD, Stone P (2013) Auction-based autonomous intersection management. In: Int. IEEE conf. on intelligent transportation systems. IEEE, pp 529–534
12. Chen W, Guha RK, Kwon TJ, Lee J, Hsu Y-Y (2011) A survey and challenges in routing and data dissemination in vehicular ad hoc networks. Wirel Commun Mob Comput 11(7):787–795
13. Chou S-Y, Lin S-W, Li C-C (2008) Dynamic parking negotiation and guidance using an agent-based platform. Expert Syst Appl 35(3):805–817
14. Deng Y, Chen Z, Yao X, Hassan S, Wu J (2019) Task scheduling for smart city applications based on multi-server mobile edge computing. IEEE Access 7:14410–14421
15. Dresner K, Stone P (2008) A multiagent approach to autonomous intersection management. Journal of Artificial Intelligence Research 31:591–656

16. HajiRassouliha A, Taberner AJ, Nash MP, Nielsen PMF (2018) Suitability of recent hardware accelerators (dsps, fpgas, and gpus) for computer vision and image processing algorithms. Signal Processing: Image Communication 68:101–119
17. Mariani S, Cabri G, Zambonelli F (2020) Coordination of autonomous vehicles: taxonomy and survey. arXiv:2001.02443
18. Melis A, Mirri S, Prandi C, Prandini M, Salomoni P, Callegati F (2016) Crowdsensing for smart mobility through a service-oriented architecture. In: 2016 IEEE int. smart cities conference (ISC2). IEEE, pp 1–2
19. Murthy DK, Masrur A (2016) Braking in close following platoons: the law of the weakest. In: 2016 Euromicro conference on digital system design (DSD), pp 613–620
20. Pinciroli C, Beltrame G (2016) Swarm-oriented programming of distributed robot networks. Computer 49(12):32–41
21. Revathi G, Dhulipala VRS (2012) Smart parking systems and sensors: a survey. In: Proc. international conference on computing, communication and applications, pp 1–5
22. Rios-Torres J, Malikopoulos AA (2017) A survey on the coordination of connected and automated vehicles at intersections and merging at highway on-ramps. IEEE Trans Intell Transport Syst 18(5):1066–1077
23. Schepperle H, Böhm K (2007) Agent-based traffic control using auctions. In: Int. workshop on cooperative information agents. Springer, pp 119–133
24. Vasirani M, Ossowski S (2012) A market-inspired approach for intersection management in urban road traffic networks. J Artificial Intell Res 43:621–659
25. Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. The Journal of Finance 16(1):8–37