Degrees of Autonomy in Coordinating Collectives of Self-Driving Vehicles

Stefano Mariani and Franco Zambonelli(✉)

Department of Sciences and Methods of Engineering,
University of Modena and Reggio Emilia, Via Amendola 2, 42122 Reggio Emilia, Italy
{`stefano.mariani,franco.zambonelli`}`@unimore.it`

Abstract. Our streets will be soon populated by multitudes of self-driving vehicles, calling for appropriate solutions to coordinate their collective movements in order to ensure safety and efficiency. In this paper, after introducing the general issues associated to coordination of self-driving vehicles, we show that a key engineering issue is identifying the most suitable degree of autonomy in decision making that should be left to vehicles during the coordination process. This issue also includes the possibility, depending on factors such as traffic conditions or the need to enact specific mobility policies, to dynamically adjust such degree of autonomy and thus the adopted coordination scheme. This introduces many theoretical and practical challenges in modelling self-driving vehicles coordination schemes and in their rigorous engineering, as in the case of intersection crossing, analysed in the paper.

Keywords: Self-driving vehicles · Coordination · Autonomy · Intersection crossing

1 Introduction

Autonomous self-driving vehicles will soon populate our streets [8]. Besides the advantage of relieving us from the duty of driving and paying attention, thus making it possible to exploit travel time in other activities, self-driving vehicles will bring further important benefits. They will reduce crashes, now mostly due to bad human behaviours and human errors, most likely saving millions of injuries and lives. They will notably reduce the number of circulating vehicles and, also thanks to route optimisation, will definitely reduce traffic and pollution [22]. Moreover, they will pave the way for a number of innovative solutions in the provisioning of mobility services [30], to serve our needs with much greater levels of quality and efficiency than today: car sharing, where fleets of autonomous vehicles (whether provided by public actors or by private companies) will be available to serve our urban mobility needs; personalised public transport and ride sharing, where autonomous vehicles and buses can dynamically gather people based on their actual required routes; smart and more effective parking approaches, in that autonomous vehicles can search for parking

T. Margaria and B. Steffen (Eds.): ISoLA 2020, LNCS 12477, pp. 189–204, 2020.
https://doi.org/10.1007/978-3-030-61470-6_12

slots based on criteria different from the "very soon and very close" one that we human usually adopt, and exploiting additional information that they might have.

The current focus of industrial and applied research in the area is on the methods and tools to enable *individual* vehicles to hit the road safely, there including rigorous engineering approaches based on formal specification and verification languages and models [34]. However, to get full advantage of self-driving vehicles, a number of situations will compulsory require *coordinating* the relative activities and movements of vehicles [1,31]. Examples of very diverse situations that require a proper and careful coordination amongst collectives of vehicles include: crossing intersections, entering a motorway, platooning, organising urban deployment and rides for fleets of ride/car-sharing vehicles, trying to improve parking occupancies and reduce parking times. Effectively supporting such coordination implies devising effective mechanisms and strategies to support coordination activities.

In this paper, we overview the key issues associated to the coordination of autonomous self-driving vehicles, and the possible approaches to attack the problem, with a specific focus on the problem of intersection crossing. We show that the different approaches to the problem are characterised by different *degree of autonomy in decision making* that is left to individual vehicles during the coordination process (not to be confused with the "level of autonomy" in driving as defined in [14]). Following, we argue that a single approach can hardly suit all possible situations. On this base, we also introduce the concept of *adjustable autonomy* in coordination, and discuss the many issues associated with the rigorous engineering of the behaviour of self-driving collectives based on adjustable autonomy. Finally, we also discuss additional general research challenges in the engineering of collectives of self-driving vehicles.

The paper is organised as follows. Section 2 characterises the problem of coordinating collectives of self-driving vehicles according to a taxonomy of coordination problems and analyses the spectrum of possible solutions; Sect. 3 focusses on the specific case of intersection crossing; Sect. 4 introduces and analyses the concept of adjustable autonomy in coordination; Sect. 5 introduces additional research challenges and Sect. 6 concludes the paper.

2 Coordinating Collectives of Self-Driving Vehicles

The problem of coordinating a collective of self-driving vehicles can be modelled as a *decision making process*, involving vehicles themselves and possibly some infrastructural entities, aimed at orchestrating vehicles' actions so as to achieve a *goal* which cannot be achieved, or not optimally, by each vehicle in isolation.

Depending on the specific nature of the coordination problem, the goal may be: *(i)* shared among a collective of vehicles, such as allocating parking slots to the vehicles of a company fleet, making the coordination problem *collaborative*, or *(ii)* individual, and possibly contrasting with the ones of the other vehicles of the collective, such as finding a parking slot in the presence of multiple private cars looking for a slot in the same zone, making the coordination problem *competitive*.

Independently of the specific problem, the issues to be faced by vehicles in coordinating are associated to the fact that in order to achieve the coordination goal, whether individual or shared, vehicles may have to: *(i)* acquire access to a shared limited resource, such as a shared intersection, calling for the coordination process to safely regulate such access according to specific strategies and rules; *(ii)* completing a specific task, such as bring a group of persons home, calling for the coordination process to properly allocate to vehicles the responsibilities and actions required to achieve such task.

The above characterisation of the problem of coordinating collectives of self-driving vehicles fits well the general characterisation of coordination problems in the areas of distributed systems [10] and multiagent systems [25]. Likewise, coordination has to satisfy the following general properties, often subject to formal verification: *safety*, expressing that "something bad never happens" during the coordination process, such as that two cars never crash while crossing an intersection; *liveness*, expressing that "something good eventually happens", such as that all cars will eventually manage to cross an intersection; *quality*, that is it should solve the coordination problem in a way to optimise some specific quality measure, such as the average or cumulative delay at which cars manage to cross an intersection.

2.1 Overview of Coordination Problems

Intersection crossing is the most representative coordination problem in urban areas, definitely the most studied in the literature, and the one we adopt as case study in this article. Intersection crossing concerns the need of coordinating vehicles while concurrently crossing intersections [29]. As such, it is a *competitive* problem in which vehicles are self-interested agents willing to obtain the right-of-way as soon as possible across the shared *resource* represented by the intersection. A proper solution to intersection crossing should enable vehicles to *safely* cross an intersection by avoiding collisions, eventually giving each vehicle the right-of-way (*liveness*), while possibly minimising the average *delay* experienced by vehicles in waiting the right-of-way (*quality measure*).

Smart parking is the problem of coordinating vehicles to have access to parking resources [27]. It assumes a competitive form in the case of private vehicles and a *collaborative* form in the case of "fleets" of vehicles made available by private companies or municipalities. A solution to this problem requires safely avoiding overbooking and starvation, whereas quality concerns the timeliness and the distance at which parking is found.

Ride sharing is the problem of coordinating vehicles to collectively satisfy "mobility tasks", such as carrying people around [4]. As in the case of parking, it has a *competitive* form for privately owned vehicles and a *collaborative* one for fleets of vehicles belonging to the same owner. Here, safety concerns assigning mobility requests only once whereas liveness amounts to guaranteeing that no user is excluded. Quality may concern maximizing car usage and monetary gains, or minimise waiting time and limit the walking distance, depending on the perspective (users or owners).

Ramp merging deals with the task of entering and leaving highway ramps [29], and requires a *cooperative* solution, as if the vehicles on the main lane are selfish, those in the merging lane will starve (hence the system as a whole would fail to achieve its goal), and if the selfish vehicles are those on the merging lane, they could cause crashes for those on the main lane, or congestion by making them slow down or change lane abruptly. Avoiding collisions and starvation while minimising the time required to perform the task are the non-functional properties to care about. Lane changing is an equivalent problem.

Platooning deals with the task of coordinating manoeuvres of vehicles so that they travel altogether as a single entity, for instance by keeping the same speed and relative positions [15]. As in ramp merging, vehicles in the platoon have incentives to cooperate with each other. Safety again concerns avoiding collisions, whereas liveness and quality mostly deal with preserving the platoon while optimising measures such as fuel consumption or speed.

Traffic flow optimisation is the large-scale coordination problem subsuming all the other ones, as it is meant to achieve a balanced exploitation of road resources so as to limit traffic congestion [33]. The problem requires *cooperative* solutions, as the routing plans that vehicles elaborate together can facilitate all individual goals. Safety amounts to avoiding congestion and traffic jams, while liveness amounts at routing vehicles so as to avoid loops and never-ending trips. A quality metric could instead be the degree of balance in the exploitation of the road network (measured by a density map) and the overall fluidity of the traffic flow (measured by the throughput of selected roads).

2.2 Coordination Solutions and Decision Making Autonomy

Independently of the specific coordination problem addressed, approaches to tackle it can be classified in terms of the *degree of autonomy in decision making* left to vehicles during the coordination process.

By "degree of autonomy in decision making" we refer to the extent to which vehicles can decide their own course of actions by themselves while coordinating. Such decision of course can be based on information acquired by vehicles about the current state of the affairs, information that can be obtained by the vehicles' own sensors, by road-side infrastructural elements, or from the other vehicles participating in the coordination process.

By definition, in any coordination process, the entities to be coordinated cannot act completely freely, and must undertake actions that account for the actions of the other entities involved in the process [10]. Thus, there is never full autonomy and freedom. However, different approaches to coordination may leave to entities different degrees of freedom in their decision making, i.e., in selecting the actions to perform during the process. Hence, the degree of autonomy may range from fully externally imposed actions (lowest autonomy) to fully self-determined action (highest autonomy), with a direct impact on the difficulty

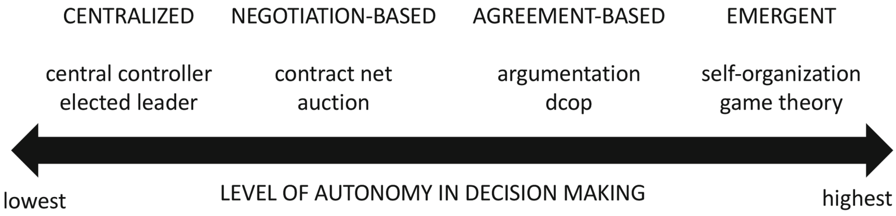

Fig. 1. Coordination approaches and level of autonomy in decision making.

of applying rigorous engineering techniques such as formal specification and verification. In particular, as depicted in Fig. 1, we can identify four main classes of coordination approaches centred around the concept of degree of autonomy: centralised, negotiation-based, agreement-based, and emergent.

In *Centralised* approaches the burden of coordination, that is, the decision making determining the outcome of the coordination process, is entirely charged upon an individual computational entity, i.e., a *coordinator*, whose decisions on how everyone should act are undebatable, and to whom vehicles must abide by design without any autonomous decision making left to them. A traditional traffic light exemplifies the role of such centralised coordinator. We emphasise that the term "centralised" here refers to the decision making process, not to the actual computing infrastructure supporting it, which can include for instance distributed processing of information by multiple sensors/cameras and/or services to perform reasoning in the Cloud.

In *Negotiation-based* approaches the burden of coordination is distributed amongst the ensemble of coordinating vehicles, who participate in a specific *negotiation protocol*, typically inspired by economic mechanisms. In a negotiation protocol, the vehicles involved can "propose" solutions and actions, each according to its own internal strategy and its own situation and goals, amongst a set of admissible moves dictated by the protocol at each step. If properly designed, the protocol will eventually guarantee the convergence towards an equilibrium solution, determining *who* (i.e., which vehicle) should do *what*, and *when*, to solve the coordination problem. Most representative negotiation protocols are: *Contract Net*, for collaborative problems, and *auctions*, for competitive ones [18].

In *Agreement-based* approaches vehicles participate in a *dynamic protocol* defined by themselves in a collective way, in a sort of dynamic meta-coordination process whose outcome is both the set of admissible moves, now jointly defined, and possibly even a re-determination of the goals to be achieved during the coordination process. The distinguishing feature here is the ability of agents to collectively define the protocol itself, that is, the goal to pursue and their strategy to make moves. Examples of these dynamic protocols include those based on *argumentation* [28], where involved entities discuss and argue together to reach a common perspective on situations, goals, and solutions, and *distributed* *constraint optimisation* [19], where agents try to collectively find a solution to an optimisation problem.

In *Emergent* approaches vehicles do not explicitly engage in any coordination protocol, thus do not even share the goal of reaching a common agreement. Rather, every vehicle behaves in a selfish way according to its goals and to maximise utility of actions w.r.t. the goals, and according to the perceptions it collects about other participants to the coordination process. It is worth emphasising that this does not contrast with the achievement of a systemic, shared goal: for instance, in ant colonies, individual ants pursue their own goal of travelling between the nest and the food source as quickly as possible, but depositing pheromone while doing so (an innate behaviour, not a coordination act) delivers the systemic goal of finding the shortest path despite disruption (e.g. due to adverse weather). Examples include: *game theoretic* approaches [24], where explicit communication is lacking, each vehicle merely assumes rationality of others, and computes its own course of actions based on informed guesses about others' expected behaviour; and *self-organising algorithms* [20], typically nature-inspired, where vehicles act in a purely reactive way, based on the implicitly perceived presence and state of other vehicles, typically expressed via "traces" in the environment, such as virtual pheromones or virtual computational fields.

3 The Case of Intersection Crossing

Today, intersection crossing is managed either by a central controller, the traffic light, or by imposing to vehicles (i.e., to their drivers) pre-defined coordination rules to be obeyed, such as stop at sign or give right-of-way to vehicles coming from the right.

In the future, thanks to self-driving vehicles, it will be possible to conceive a variety of innovative solutions, safer and more efficient, eventually making traffic lights and stop signs obsolete. Based on the classification of coordination solutions along the level of autonomy in decision making, let us now overview the variety of such solutions.

3.1 Centralised

Centralised approaches to intersection crossing assume the existence of a computational central authority, the intersection manager, bearing alone the burden of decision making. It is typically in charge of: *(i)* receiving information from vehicles approaching the intersection (i.e., origin, destination, speed); *(ii)* elaborating a set of collision free trajectories enabling vehicles to safely cross the intersection, which may require some vehicles to slow down or change lane; and *(iii)* instructing, or directly commanding, the vehicles about what to do, or informing them about what constraints they must abide to while crossing the intersection. Centralised approaches are usually the easiest to rigorously engineer, as the whole coordination algorithm is executed by a single component amenable of formal verification, without vehicles autonomy to hinder the process.

Examples of centralised proposals to intersection crossing include [40], which attack the problem in terms of a traditional mutual exclusion approach, and [16], in which the authors propose a control algorithm implementing a nonlinear constrained optimisation in charge of computing the best moves for every vehicle and then directly manipulating vehicles' driving parameters. A similar stance is taken in [41], where cooperative adaptive cruise control is exploited for intersection crossing by assuming that a smart controller device placed in the intersection can communicate with incoming vehicles to instruct them about the actions to perform.

Other approaches are a little more permissive and let the inbound vehicles decide how to fulfil a set of constraints set by the intersection manager, which may regard the time slot assigned for crossing, as in the work by Dresner and Stone [7]: the authors propose a *reservation-based approach* in which incoming vehicles request assignment of a time slot for crossing to the intersection manager, who computes decisions based on a local control policy.

In general terms, all the above approaches ensure safety and avoid starvation by giving every vehicles the possibility to cross the intersection. Most importantly, simulations show that such approaches dramatically reduce the waiting time for vehicles with respect to traditional approaches based on stop signs or traffic lights [7], because: *(i)* the occupancy of the intersection is maximised and *(ii)* vehicles from different directions can cross the intersection without waiting, provided they are not in direct collision—i.e., they occupy different portions of the intersection, or occupy the same portion at different times.

A problem of centralised approaches is requiring the presence of a dedicated infrastructural element (the intersection manager) and the capability of vehicles to interact with it at all times. Thus, they can hardly be applicable in the wild. Also, such a central authority is an obvious bottleneck for both performance and tolerance to failures. A recent proposal [36] suggests the possibility for vehicles to be engaged in a leader election algorithm, to elect a transient leader vehicle in charge to act as intersection manager for a predefined amount of time.

3.2 Negotiation-Based

In negotiation-based approaches, vehicles are required to actively participate in a protocol aimed at establishing in which order vehicles will gain access to the intersection. Such protocol, given the competitive nature of the problem, can take the form of an *auction*. In approaching the intersection, vehicles may contact an intersection manager or a broker temporarily elected amongst themselves, by placing a "bid", that is, by making an offer to "buy" the portions of the intersection they require for crossing, for the time required to cross. The value of the bid expresses the urgency of the vehicle in crossing, it is autonomously set by each individual vehicle according to its own strategy, and can correspond to some real-world currency or some sort of "road credits" assigned to vehicles. The broker collects the bids, gives the right-of-way to the set of vehicles that are in a collision free trajectory and, amongst those that are in collision, to the ones having placed the highest bid.

Examples of auction-based protocols for intersection crossing are described in [38], [6], and [5]. There, different policies to resolve the auction are analysed, based on different strategies put in place by the bidding vehicles, as well as different strategies by the broker in establishing the winners. Such strategies can also attempt at incentivising fair bidding while discouraging malicious behaviours.

Auction-based mechanisms, with slight variations depending on the adopted strategies, generally exhibit performances comparable and at times superior to that of centralised ones: the waiting time of vehicles is dramatically lower than that of traditional traffic lights. Safety is ensured provided that vehicles respect the "rules of the game", and accept waiting when losing the auction. A problem intrinsic to any auction mechanism concerns liveness, that is, the property of having every vehicle achieve its goal without starvation, in that the strategy of bidding vehicles can sometimes make others to experience indefinitely long waiting times. Also, if implemented through a dedicated intersection manager acting as broker, they inherit the "bottleneck" drawbacks of centralised approaches.

3.3 Agreement-Based

Intersection crossing with agreement techniques essentially amounts to give vehicles approaching an intersection the possibility to interact so as to affect each others' original goals (e.g., directions) and priorities.

An example proposal, specifically conceived in the context of bimodal traffic (vehicular plus public transport), is discussed in [3]. There, agreement between vehicles happens through a repeated communication protocol running between approaching vehicles and buses, with the assistance of an heterogeneous pool of agents representing conflicting goals, such as the need to minimise private vehicles travel time while prioritising public transportation. Depending on both macro and micro scale criteria, in fact, the agents participating in the protocol may decide to prioritise, hence, ultimately, giving right-of-way to, either private traffic or public transportation, as a result of a conflict resolution protocol.

Another interesting approach models intersection crossing as a *Distributed Constrained Optimisation Problem* (DCOP) [39], that is, interpreting vehicles as a multi-agent system in which agents have to find an agreement about the best solution possible to a dynamic set of shared constraints. This kind of modelling lends itself to a distributed implementation, where each vehicle interacts in a peer-to-peer way with the neighbouring ones to solve a local problem, that is, DCOP limited to those vehicles actually approaching the intersection. For doing so, the involved agents actually resort to a messaging protocol to exchange their current solutions as they try to adjust their individual values to converge to a feasible (hopefully, optimal) solution.

Finally, let us mention the approach we envisioned in [17], which proposes to adopt *argumentation technologies*. In particular, we suggest vehicles can engage in open dialogues while approaching the intersection, discussing their beliefs about the best way to approach the intersection, and in case of conflicting needs, arguing with each other about possible ways to avoid that conflict. During the dialogue, vehicles can change the argumentation strategy, and may evaluate assertions differently based on the dynamic contingencies arising in the meantime. For instance, a vehicle A approaching the intersection in the north-to-south direction can express arguments about its urgency to cross, and can argue that another vehicle B in the east-to-west direction (and thus conflicting with A) could/should decide to cross right, as that move would make B reach destination anyways, but would avoid the conflict with A. Persuaded by solidity of A's argument, B could eventually decide to turn right. Although still at the conceptual level, an argumentation-based approach to intersection crossing shows potential for greater flexibility and adaptivity in facing unforeseen situations. In addition, the power of argumentation approaches in the area of autonomous driving is advocated also by other conceptual proposals as a way to solve conflicts and increase trustworthiness and safety of decisions [9].

3.4 Emergent

Handling intersection crossing with coordination by emergence implies giving absolute freedom to vehicles in choosing how to cross intersections, with the only constraints of acting in a safe way and avoiding starvation. To this end, one can let vehicles either: *(a)* play a selfish game where each agent attempts to maximise its expected utility in crossing despite other agents' needs and goals [24]; or *(b)* be engaged in an implicit, self-organising coordination scheme, where each vehicle responds in a reactive way to the actions of the other agents, according to some sort of "natural laws" enforced in the intersection "ecosystem" [26]. In both cases, coordination does not consider an explicit agreement about what to do.

In [21], the authors interpret the intersection crossing problem using *game theory*, that is, modelling each vehicle as the player of a game involving other approaching vehicles, each playing its own game, thus each having different payoffs and utility models—that is, essentially, each player is unaware of the formalisation of the game others are playing. The proposed approach investigates how to build decision matrices in such a way that minimal information can be assumed by agents while still being able to find a solution for their own game—that is, a safe way to cross the intersection in the lowest possible time. Alternatively, it is possible to model the collective behaviour of vehicles at intersections in terms of a self-organised collective movement, similar to that of flocking birds [35].

Actually deploying autonomous vehicles that cross intersections by relying on such approaches seems hardly feasible, as delivering guarantees about safety and liveness may be prohibitively difficult or impractical in the general case of emergent approaches to coordination, because these approaches often exploit stochastic decision making and partial, local information. Possibly, however, in mixed scenarios with the presence (as discussed in Subsect. 5.3) of non connected human-driven vehicles, emergent approaches can be the only solutions for individual vehicles to coordinate with each other.

4 Adjustable Autonomy

In previous section, for the specific coordination problem of intersection crossing, we have presented different approaches based on a different degree of autonomy left to vehicles. The selection of the best strategy, though, may depend on the specific current traffic situation at an intersection, and a single solution can hardly handle all possible situations optimally. For instance, in the case of intersection crossing:

- A solution based on distributed negotiation or argumentation between vehicles can be very effective in rather low traffic situations, when the number of vehicles involved in such negotiations is quite low, thus a collective outcome can be reached quickly because the number of messages to exchange even in the case of completely connected topology would remain low.
- In the case of congested traffic situations, with a large number of vehicles involved, reaching a shared agreement can be harder and induce notable overhead and delay in communications. Also, in the case of auctions, it can induce inflationary effects on the bids. In these situations, thus, on the one hand it could be more appealing to rely on centralised solutions so as to reduce the complexity of communications (e.g. bandwidth consumption), on the other hand emergent ones may further help avoiding the bottleneck of having a single point of failure while still keeping communication costs low.
- Emergent approaches can possibly work both in very low-traffic situations and in highly congested ones, as mentioned above, for their capacity to scale seemingly with the scale of the problem (e.g. as regards communication costs and computational complexity of the protocol), but further experiments are needed to confirm this opportunity.
- Likely, as discussed in [13], it can be necessary for different intersections to adopt different coordination schemes at different times in order to support traffic flow optimisation.

Similar issues can apply also to other classes of coordination problems. For example, consider the need to coordinate vehicles in order to optimise the usage of parking slots. In general, a centralised parking scheme that works well to let the city governance control the distribution of parked vehicles, may fall short in the presence of a high number of vehicles by inducing notable delays in parking. In this case it is better to switch to an approach that lets individual vehicles negotiate for parking slots according to their own preferences.

All the above considerations suggest the possibility, for coordinating vehicles to properly enforce safety, liveness, and to maximise quality in coordination, to dynamically switch from one coordination scheme to another upon changing conditions. Indeed, many municipalities already often adopt a similar dynamic adaptation of the scheme to regulate intersection crossing: the traffic lights that regulate access to an intersection during the day (i.e., in situations of expected intense traffic) are switched off at night (expected low traffic) to let vehicles directly coordinate with each other.

In the area of robotics and multiagent systems the theme of “coordination with adjustable autonomy” [23] (sometimes referred to as “flexible autonomy”, also [11]) has been extensively discussed, either referring to the fact that, at times, a human actor may wish to reclaim autonomy in decisions from agents or robots [32], or to the fact that (as in our scenario) specific conditions may require to dynamically switch the coordination scheme [37].

In the real world, and in the context of safety-critical situations such as those involving the coordination of autonomous vehicles, though, designing and realising such dynamic switch in a rigorous and reliable way can be conceptually and technically very hard, and requires facing several challenges. In particular:

- For evaluating the switch to a different coordination scheme, there is the need of well-defined metrics and background knowledge to evaluate which situation fits which coordination scheme. This requires extensive simulations and real-world experiences to compare the effectiveness of different schemes in different situations. For instance, based on the description of the different autonomy classes and their representative protocols, it is likely that factors to consider while deciding which scheme to adopt include *(i) raw performance* aspects, such as the number of messages exchanged and the number of iterations the protocol needs to converge (both tend to increase with decentralisation), as well as *(ii)* accounting for the *amount of information* needed for the protocol to work, such as whether it needs global information (for which a centralised approach may be the only reliable option) or not, and finally *(iii) liability* issues, to establish individual responsibilities in case something bad happens (the more autonomy vehicles have, the less an individual responsible is likely to be found). Whether for performance we might already have the right tools to rigorously measure it, the same does not hold for the information and liability aspects, which remain open issues.
- Identifying the situation for a switch requires continuous detailed monitoring of the traffic situation and of the effectiveness of the coordination process. Also, predictive monitoring techniques should be adopted, to let the switch take place before a degradation of quality in coordination occurs. This obviously implies having means to precisely measure coordination effectiveness, which as far as we know are not widely established, yet.
- Deciding the actual process by which the switch should be decided and enacted. In other words, the vehicles and/or the centralised manager involved in coordination should agree on the switch and on the actions by which to actually perform it. That is, there must be a meta-coordination protocol taking place for the switch to a different coordination scheme. This could again rely on a centralised controller to decide and enact it, or on vehicles negotiating/agreeing with each other on how and when to switch, as it may be required in the absence of infrastructures supporting the existence of a centralised controller. Hence, the same considerations we made for the coordination protocols regarding performance and autonomy trade-offs, will apply to this meta-coordination layer, too, further complicating the open challenge of deciding when to switch.

All of the above, should lead to solutions with provable properties of stability and with provable convergence times. We do not have solutions ready to use to propose here, but certainly the vast amount of literature on adjustable autonomy can suggest useful research directions.

5 Additional Research Challenges

Let us now introduce a few additional general challenges, i.e., issues concerning vehicles coordination beside the specific case of intersection crossing. Autonomous vehicles can hardly be deployed in the real-world and start coordinating without also identifying rigorous solutions to these challenges.

5.1 Systemic Coordination

So far, we discussed the issue of coordination mostly at the level of individual, isolated systems, such as a single intersection. However, thinking at a more systemic level, such as at urban scale, coordination actions in one part of the system may indeed impact other parts of the system. For example, queues at an intersection can induce queues at nearby intersections, or a slow down in a motorway due to an intense flux of traffic in an entering lane can quickly propagate backwards to impact previous entering lanes.

The inter-related effects of individual coordination acts along with the need to respect global level policies imply that the solutions and the policies adopted to solve an individual coordination problem cannot be designed without accounting for the systemic impact of such solutions and policies. In other words, the level of individual coordination must be coupled with a co-coordination one, in which an agreement at the global level is reached on how to act, that is, according to which policies and constraints, at the local level. In the area of autonomous vehicles, a few works exist that handle such systemic problems. For instance, [38] analyses how global coordination of intersections can be achieved by trying to affect, at the local level, the choices of individual vehicles. A similar analysis is presented in [13]. In different fields such as logistics, energy management, robotics, and multi-agent systems, a variety of mechanisms have been proposed for coordination in large-scale systems of systems: hierarchical mechanisms, market-based, self-organising [25]. Such mechanisms can be a source of inspiration for the field of autonomous vehicles as well, but would also call for tools to enable accurate simulation and prediction of the global impacts of coordination solutions.

5.2 Intersection Markets

Today, while driving, we are already used to pay for the usage of infrastructures such as parking slots, bridges, motorways. However, these payments are based on static pricing schemes and offer a neutral service. If, as in negotiation-based solutions, vehicles can dynamically request access to intersections, or to other road infrastructures such as parkings and motorways, and pay them automatically, it may become possible for the manager of such infrastructures to impose dynamic pricing mechanisms, based on the current demand. Doing so implies that a vehicle, while starting its ride, may have no a priori idea about how much it will eventually cost. Also, this opens up the way for imposing fees on intersection crossing, imposing payments for crossing busy intersections with fees varying depending on traffic and time-to-wait.

The mechanisms of dynamic payments could also enable a model in which passengers can decide to pay more to get better services, e.g., crossing an intersection quickly, breaking the current neutrality of road infrastructures. In the future, such mechanisms could become based on a real auction with real money, with the consequence that vehicles whose owners/drivers/passengers have higher budget will always bid higher and buy priority in crossing the intersection, while vehicles whose owners/drivers/passengers have lower budget will risk starvation. The above issues, other than calling for proper algorithmic solutions to avoid unfairness or inflationary effects, also call for the definition of suitable regulations to avoid mobility becoming a privilege, and suitable means to integrate such regulation into a coordination scheme.

5.3 Mixed Scenarios

This article assumes that all the cars are fully autonomous, or at least that they act and interact autonomously with each other during the coordination act. This can match a not-so-near future when we can expect that human-driven cars will no longer exist or when, for safety and efficiency reasons, it will be forbidden for humans to drive but in specific controlled situations: the same as today, for instance, it is forbidden to ride a horse in motorways and high-speed roads. However, there will be a rather long transition phase in which our streets will be populated by a mixture of fully autonomous cars, partially autonomous ones, and traditional human-driven cars, other than bicycles and motorbikes. Such a scenario clearly challenges the possibility of relying on the surveyed coordination schemes, unless one devises dependable means to involve human-driven cars (that is, their human drivers) in the process of coordination, and suitable means to formalise and ensure properties in such mixed scenarios.

These scenarios have several characteristics in common with the issue of coordinating the movements of mixed teams of robots and humans, which has been extensively analysed in the context of robotics and autonomous systems [12], but assumes the existence of means for robots and humans to communicate with each other. Also, designing a coordination scheme according to the solutions discussed in this article would require accounting for the possible inaccuracy of actions by human-driven cars, and possibly for the presence of non-connected vehicles (e.g. bicycles) whose behaviour cannot be rigorously predicted. For this latter case, emergent approaches to coordination could provide solutions, though possibly much less effective.

6 Conclusions

For future autonomous vehicles to populate our streets, it will be necessary to identify rigorous solutions for coordinating their relative movements in order to let them circulate safely and without conflicts and crashes. In this article, we focussed on the problem of crossing intersections, and showed how a variety of solutions can be conceived, each characterised by a different level of autonomy in decision making left to vehicles during coordination.

Selection of the appropriate solution to handle intersection crossing will require proper modelling of the problem and of the domain, other than rigorous approaches to analyse and compare the different solutions, in order to select the most appropriate one depending on the context. Furthermore, it will require addressing a number of additional challenges that represent promising directions for future research.

Finally, but this problem would require a full analysis on its own, the safety-critical nature of autonomous vehicles will possibly require them to solve ethical dilemmas while coordinating with each other, e.g., multi-vehicle instantiations of the trolley problem [2], which raises the additional issue of somewhat engineering a sort of moral dimension for vehicles.

References

1. Abeywickrama, D.B., Mamei, M., Zambonelli, F.: Engineering collectives of self-driving vehicles: the SOTA approach. In: Margaria, T., Steffen, B. (eds.) ISoLA 2018. LNCS, vol. 11246, pp. 79–93. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03424-5_6
2. Awad, E., et al.: The moral machine experiment. Nature **563**(7729), 59 (2018)
3. Balbo, F., Bhouri, N., Pinson, S.: Bimodal traffic regulation system: a multi-agent approach. Web Intell. **14**(2), 139–151 (2016). https://doi.org/10.3233/WEB-160336
4. Bicocchi, N., Mamei, M., Sassi, A., Zambonelli, F.: On recommending opportunistic rides. IEEE Trans. Intell. Transp. Syst. **18**(12), 3328–3338 (2017). https://doi.org/10.1109/TITS.2017.2684625
5. Cabri, G., Gherardini, L., Montangero, M.: Auction-based crossings management. In: Proceedings of the 5th EAI International Conference on Smart Objects and Technologies for Social Good, GoodTechs 2019, pp. 183–188. ACM, New York (2019). https://doi.org/10.1145/3342428.3342689
6. Carlino, D., Boyles, S.D., Stone, P.: Auction-based autonomous intersection management. In: 2013 16th International IEEE Conference on Intelligent Transportation Systems-(ITSC), pp. 529–534. IEEE (2013)
7. Dresner, K., Stone, P.: A multiagent approach to autonomous intersection management. J. Artif. Intell. Res. **31**, 591–656 (2008)
8. Fagnant, D.J., Kockelman, K.: Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations. Transp. Res. Part A Policy Pract. **77**, 167–181 (2015)
9. Fridman, L., Ding, L., Jenik, B., Reimer, B.: Arguing machines: human supervision of black box AI systems that make life-critical decisions (2017)

10. Gelernter, D., Carriero, N.: Coordination languages and their significance. Commun. ACM **35**(2), 96–107 (1992). https://doi.org/10.1145/129630.376083
11. Gerber, C., Siekmann, J., Vierke, G.: Flexible autonomy in holonic agent systems. In: Proceedings of the 1999 AAAI Spring Symposium on Agents with Adjustable Autonomy. AAAI (1999)
12. Goodrich, M.A., Schultz, A.C., et al.: Human-robot interaction: a survey. Found. Trends Hum. Comput. Interact. **1**(3), 203–275 (2008)
13. Hausknecht, M., Au, T.C., Stone, P.: Autonomous intersection management: multi-intersection optimization. In: 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 4581–4586. IEEE (2011)
14. International, S.: Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles, June 2018
15. Jia, D., Lu, K., Wang, J., Zhang, X., Shen, X.: A survey on platoon-based vehicular cyber-physical systems. IEEE Commun. Surv. Tutorials **18**(1), 263–284 (2016). (Firstquarter). https://doi.org/10.1109/COMST.2015.2410831
16. Lee, J., Park, B.: Development and evaluation of a cooperative vehicle intersection control algorithm under the connected vehicles environment. IEEE Trans. Intell. Transp. Syst. **13**(1), 81–90 (2012). https://doi.org/10.1109/TITS.2011.2178836
17. Lippi, M., Mamei, M., Mariani, S., Zambonelli, F.: An argumentation-based perspective over the social iot. IEEE Internet Things J. **5**(4), 2537–2547 (2018). https://doi.org/10.1109/JIOT.2017.2775047
18. Lopes, F., Wooldridge, M., Novais, A.Q.: Negotiation among autonomous computational agents: principles, analysis and challenges. Artif. Intell. Rev. **29**(1), 1–44 (2008)
19. Maheswaran, R.T., Tambe, M., Bowring, E., Pearce, J.P., Varakantham, P.: Taking DCOP to the real world: efficient complete solutions for distributed multi-event scheduling. In: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1, pp. 310–317. IEEE Computer Society (2004)
20. Mamei, M., Menezes, R., Tolksdorf, R., Zambonelli, F.: Case studies for self-organization in computer science. J. Syst. Architect. **52**(8–9), 443–460 (2006)
21. Mandiau, R., Champion, A., Auberlet, J.M., Espié, S., Kolski, C.: Behaviour based on decision matrices for a coordination between agents in a urban traffic simulation. Appl. Intell. **28**(2), 121–138 (2008). https://doi.org/10.1007/s10489-007-0045-3
22. Menon, N., Barbour, N., Zhang, Y., Pinjari, A.R., Mannering, F.: Shared autonomous vehicles and their potential impacts on household vehicle ownership: an exploratory empirical assessment. Int. J. Sustain. Transp. **13**, 1–12 (2018)
23. Mostafa, S.A., Ahmad, M.S., Mustapha, A.: Adjustable autonomy: a systematic literature review. Artif. Intell. Rev. **51**(2), 149–186 (2017). https://doi.org/10.1007/s10462-017-9560-8
24. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press, Cambridge (2007)
25. Omicini, A., Zambonelli, F.: Challenges of decentralized coordination in large-scale ubicomp systems. In: Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp Adjunct 2016, Heidelberg, Germany, 12–16 September, 2016, pp. 1315–1320. ACM Press (2016)
26. Parunak, H.V.D., Brueckner, S., Sauter, J.: Digital pheromone mechanisms for coordination of unmanned vehicles. In: Castelfranchi, C., Johnson, W.L. (eds.) 1st International Joint Conference on Autonomous Agents and Multiagent systems, vol. 1, pp. 449–450. ACM, New York, 15–19 July 2002 (2012)

27. Polycarpou, E., Lambrinos, L., Protopapadakis, E.: Smart parking solutions for urban areas. In: 2013 IEEE 14th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM). IEEE (2013). https://doi.org/10.1109/WoWMoM.2013.6583499
28. Rahwan, I., Ramchurn, S.D., Jennings, N.R., Mcburney, P., Parsons, S., Sonenberg, L.: Argumentation-based negotiation. Knowl. Eng. Rev. **18**(4), 343–375 (2003)
29. Rios-Torres, J., Malikopoulos, A.: A survey on the coordination of connected and automated vehicles at intersections and merging at highway on-ramps. IEEE Trans. Intell. Transp. Syst. **18**(5), 1066–1077 (2017). https://doi.org/10.1109/TITS.2016.2600504
30. Röth, T., Pielen, M., Wolff, K., Lüdiger, T.: Urban vehicle concepts for the shared mobility. ATZ Worldwide **120**(1), 18–23 (2018). https://doi.org/10.1007/s38311-017-0163-4
31. Sassi, A., Zambonelli, F.: Coordination infrastructures for future smart social mobility services. IEEE Intell. Syst. **29**(5), 78–82 (2014). https://doi.org/10.1109/MIS.2014.81
32. Scerri, P., Pynadath, D., Tambe, M.: Adjustable autonomy in real-world multi-agent environments. In: Proceedings of the Fifth International Conference on Autonomous Agents, pp. 300–307. ACM Press (2001)
33. Seredynski, M., Bouvry, P.: A survey of vehicular-based cooperative traffic information systems. In: 2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC), pp. 163–168. IEEE, October 2011. https://doi.org/10.1109/ITSC.2011.6083055
34. Seshia, S.A., Sadigh, D., Sastry, S.S.: Formal methods for semi-autonomous driving. In: 2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC), pp. 1–5 (2015)
35. Toner, J., Tu, Y.: Flocks, herds, and schools: a quantitative theory of flocking. Phys. Rev. E **58**(4), 4828 (1998)
36. Tonguz, O.K.: Red light, green light–no light: Tomorrow's communicative cars could take turns at intersections. IEEE Spectr. **55**(10), 24–29 (2018). https://doi.org/10.1109/MSPEC.2018.8482420
37. van der Vecht, B., Dignum, F., Meyer, J.-J.C., Neef, M.: A dynamic coordination mechanism using adjustable autonomy. In: Sichman, J.S., Padget, J., Ossowski, S., Noriega, P. (eds.) COIN -2007. LNCS (LNAI), vol. 4870, pp. 83–96. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-79003-7_7
38. Vasirani, M., Ossowski, S.: A market-inspired approach for intersection management in urban road traffic networks. J. Artif. Int. Res. **43**(1), 621–659 (2012)
39. Vu, H., Aknine, S., Ramchurn, S.D.: A decentralised approach to intersection traffic management. In: IJCAI, pp. 527–533. AAAI Press (2018)
40. Wu, W., Zhang, J., Luo, A., Cao, J.: Distributed mutual exclusion algorithms for intersection traffic control. IEEE Trans. Parallel Distrib. Syst. **26**(1), 65–74 (2015). https://doi.org/10.1109/TPDS.2013.2297097
41. Zohdy, I.H., Kamalanathsharma, R.K., Rakha, H.: Intersection management for autonomous vehicles using iCACC. In: 2012 15th International IEEE Conference on Intelligent Transportation Systems, pp. 1109–1114. IEEE, September 2012. https://doi.org/10.1109/ITSC.2012.6338827