

Robert Hult, Gabriel R. Campos, Erik Steinmetz,
Lars Hammarstrand, Paolo Falcone, and Henk Wymeersch

Coordination of Cooperative Autonomous Vehicles

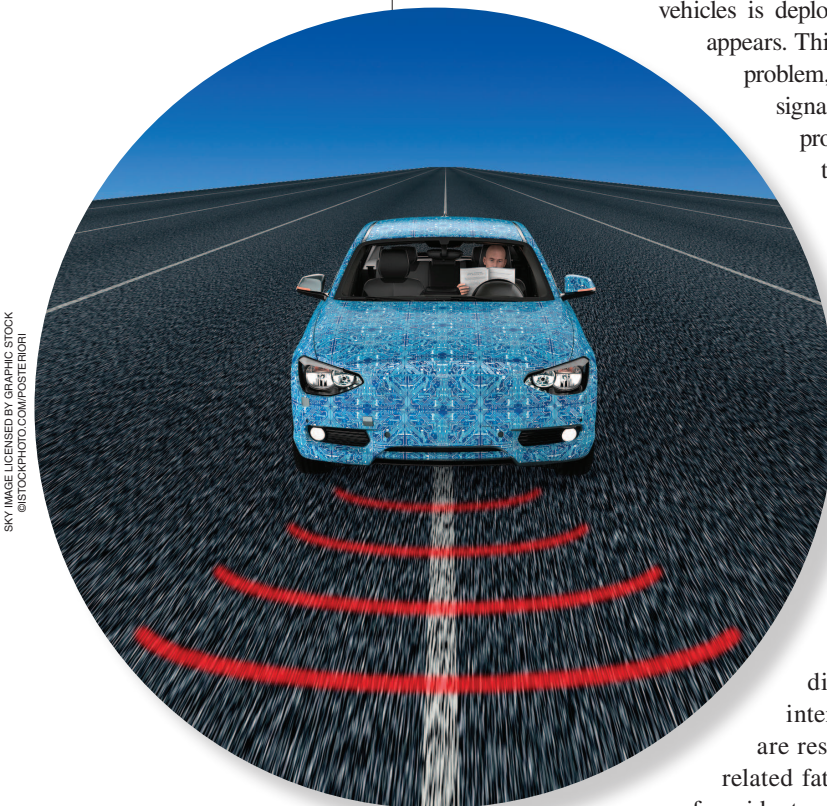
Toward safer and more efficient road transportation

While intelligent transportation systems come in many shapes and sizes, arguably the most transformational realization will be the autonomous vehicle. As such vehicles become commercially available in the coming years, first on dedicated roads and under specific conditions, and later on all public roads at all times, a phase transition will occur. Once a sufficient number of autonomous vehicles is deployed, the opportunity for explicit coordination appears. This article treats this challenging network control problem, which lies at the intersection of control theory, signal processing, and wireless communication. We provide an overview of the state of the art, while at the same time highlighting key research directions for the coming decades.

Introduction

The purpose of intelligent transportation systems (ITS) is to leverage advances in information technology to alleviate major problems in the current road traffic system. Focus areas include the prevention and mitigation of accidents, reduction of greenhouse gas emissions, and efficiency in terms of energy and infrastructure utilization.

A particularly problematic subset of traffic scenarios in terms of both safety and efficiency is those where vehicles must coordinate the use of a common resource, such as intersections, roundabouts and on ramps. These are responsible for a significant fraction of traffic-related fatalities and injuries [1]. Due to the high risk of accidents, these traffic scenarios are among the most regulated, with vehicles guided simultaneously by traffic lights, signs, road markings, and right-of-way rules. The problems of traffic fatalities and inefficiency are expected to become even more pressing in the future, as the global number of light vehicles (e.g., passenger cars and light trucks) is forecast to rapidly increase. Proportional expansion of road infrastructure is undesirable in most countries, and might not even be possible given continued urbanization and the associated increase in population density. Hence, there is great interest to



SKY IMAGE LICENSED BY GRAPHIC STOCK
©ISTOCKPHOTO.COM/POSTERION

improve safety, energy, and traffic efficiency on the existing and planned road infrastructure.

Many of the aforementioned problems are caused by the human involvement in the coordination of traffic. Studies have shown that over 90% of traffic accidents are completely, or in part, due to human error [2]. This has led to a progressive shift in responsibilities from the human driver to dedicated control systems, most recently in the form of autonomous vehicles, which aim to provide more efficient, comfortable, and virtually accident-free road traffic. Autonomous vehicles are still limited in terms of their sensing and coordination capabilities, as their actions depend on the on-board sensory data and models of other vehicles' behavior. As an example, a summary of the Urban Grand Challenge [3] mentioned that a number of incidents could have been avoided if vehicles could anticipate the behavior of other vehicles, and that vehicles should cooperate for autonomous driving to reach its full potential. The benefit of cooperation was already recognized in a parallel track in vehicle automation, and platooning, which instead of complete autonomy promotes information sharing between vehicles and joint decision making. In a platoon, the vehicles rely on vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication to share information regarding the environment and internal states and choose safe and efficient control policies jointly [4].

The two tracks for automating vehicles have thus followed different approaches: one (the platooning track) explicitly relies on communication among vehicles, while the second (the autonomous vehicle track) does not. With the adoption of the IEEE 802.11p Standard, as well as the possibilities of V2V and V2I communication and other services under the future 5G wireless standard [5], the two tracks are expected to merge, leading to a new type of large-scale wireless networked control system. This merging will likely take place in a piecemeal fashion, with first a ubiquitous availability of wireless communications, and only later the gradual introduction of cooperative autonomous vehicles [6]. These vehicles will drive autonomously, but at the same time be able to leverage their communication capabilities for cooperative planning and control, as well as cooperative perception and sensing, thus eliminating many of the traffic-safety and efficiency problems. The design and operation of such networks of cooperating vehicles place enormous demands on the control, communications, and sensing subsystems, as they must operate in harmony across different brands and types of vehicle, with limited margin for error.

In this article, we give an overview of the coupling between control, communication, and sensing, as visualized in Figure 1. We provide a survey of the different control approaches and their associated signal processing challenges. We hope that this article can provide an introduction for signal processing professionals to the control-theoretic aspects of vehicle coordination and pave the way for a tighter collaboration.

Problem formulation

The problem of coordinating a set of vehicles can be phrased as calculating the best control trajectories for the individual vehicles that allow them to safely reach their destination in finite time (e.g., within a few tens of seconds). In general, any solution should meet the basic requirements of safety (i.e., no collisions occur) and liveness (i.e., destinations are reached eventually), while optimizing some performance metric. The most important requirement is safety. Hence, vehicles may never be steered to

states from which future collisions are unavoidable. Second, the coordination algorithm must guarantee that all vehicles are allowed to both enter and exit the coordination area in finite time so that permanent stops and traffic deadlocks are avoided. Finally, a performance criterion is necessary to favor one among multiple solutions. In summary, a coordination problem can be stated as a constrained optimal control problem, where a performance criterion is optimized with respect to the vehicles' control input trajectories, subject to safety and liveness requirements:

$$\text{minimize} \quad \text{performance criterion} \quad (1a)$$

$$\text{subject to} \quad \text{safety constraints,} \quad (1b)$$

$$\text{liveness constraints.} \quad (1c)$$

The constrained optimal control framework clearly allows one to conveniently accommodate performance, safety, and

We give an overview of the coupling between control, communication, and sensing for control of automated vehicles.

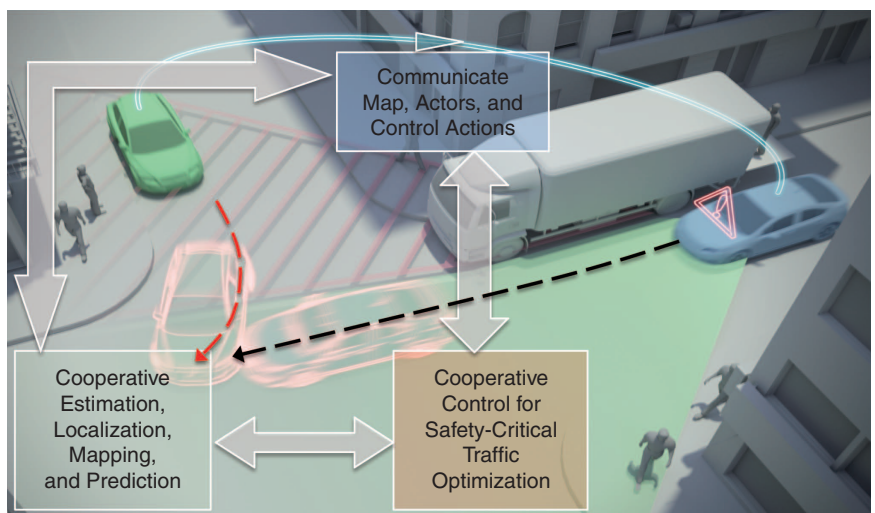


FIGURE 1. Vehicle coordination relies on a tight interaction between control, communication, and sensing.

liveness. However, as will be discussed later in this article, the partial lack of formal analysis tools limits their practical applicability. In particular, the impact of imperfect sensing data and communication impairments on stability and feasibility (i.e., the capability of finding a solution that meets safety and liveness requirements) of (1) is not completely understood under realistic communication protocols and sensing scenarios. In the absence of sensing and communication impairments, a simplified version of the generic problem (1) can be specified mathematically as follows.

Safety constraints

Consider a set of N vehicles (agents), whose motion is described by

$$\dot{x}_i(t) = f_i(x_i(t), u_i(t), t), \quad (2)$$

where $x_i \in \mathcal{X}_i \subseteq \mathbb{R}^n$ and $u_i \in \mathcal{U}_i \subseteq \mathbb{R}^m$ are the state and input/control vectors, respectively, $\dot{x}_i(t)$ denotes the time derivative of $x_i(t)$, and the sets $\mathcal{X}_i, \mathcal{U}_i$ reflect physical and design constraints. Examples of such constraints are acceleration limitations and the vehicles' minimum and maximum speeds. Examples of the state x_i are vectors comprising the vehicle's position and velocity in one, two, or three dimensions. Let $\mathcal{G}_i(x_i)$ describe the vehicle geometry, being the closed and compact set of spatial coordinates that vehicle i occupies when its state is x_i . Hence, a collision between two vehicles i and j occurs at time t if

$$\mathcal{G}_i(x_i(t)) \cap \mathcal{G}_j(x_j(t)) \neq \emptyset. \quad (3)$$

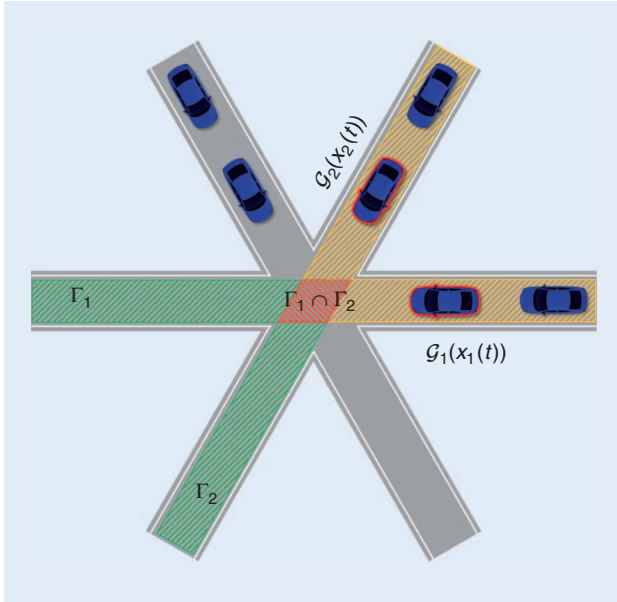


FIGURE 2. An example of a traffic coordination scenario at a three-way intersection. The geometries of vehicles 1 and 2 are highlighted in red, and their paths Γ_i are dashed and colored. The critical region, $\Gamma_1 \cap \Gamma_2$, is shown in dashed red, the target sets $\mathcal{T}_1, \mathcal{T}_2$ in green, and the paths before the critical region in yellow. The vehicle geometry $\mathcal{G}_i(x_i(t))$, depending on the vehicle state $x_i(t)$, is also depicted.

Furthermore, we denote by Γ_i the closed and connected set of spatial coordinates that comprises the paths of vehicle i so that a vehicle is on its path if $\mathcal{G}_i(x_i(t)) \subseteq \Gamma_i$. Provided that each vehicle stays on its path, a collision between vehicles can consequently only take place within a critical region where $\Gamma_i \cap \Gamma_j$, i.e., where paths fully or partly overlap (i.e., where paths are the same, cross, or merge).

Liveness constraints

Assume that all paths Γ_i are fixed and constant, and let the target set $\mathcal{T}_i \subset \Gamma_i$ be the set of spatial coordinates that vehicle i strives to reach (e.g., the road after an intersection, roundabout, or onramp). If $\mathcal{G}_i(x_i(t)) \subset \mathcal{T}_i$ is satisfied in finite time for all vehicles, the coordination is said to be deadlock-free, and all vehicles are eventually coordinated through the critical regions. For an illustration of the introduced notations, see Figure 2.

Performance criterion

In general, the cost for vehicle i , denoted by $J_i(x_i(t), u_i(t))$, can be expressed as

$$J_i(x_i(t), u_i(t)) = \int_0^{+\infty} \Lambda_i(x_i(t), u_i(t), t) dt,$$

where the stage cost $\Lambda_i(x_i(t), u_i(t), t)$ could be, e.g., instantaneous power consumption so that $J_i(x_i(t), u_i(t))$ is the total consumed energy. Other examples of $\Lambda_i(x_i(t), u_i(t), t)$ include a deviation from a target speed, or a measure of discomfort for the driver.

Overall problem and its receding horizon formulation

With the introduced notations and concepts, the N -vehicle optimal coordination problem (OCP) is now naturally formulated as the following infinite time, constrained optimal control problem.

Problem 1: OCP

$$\text{minimize}_{\mathbf{x}(t), \mathbf{u}(t)} \sum_{i=1}^N J_i(x_i(t), u_i(t)) \quad (4a)$$

$$\text{subject to } \dot{x}_i(t) = f_i(x_i(t), u_i(t), t), x_i(0) = x_{i,0} \quad (4b)$$

$$x_i(t) \in \mathcal{X}_i, u_i(t) \in \mathcal{U}_i \quad (4c)$$

$$\mathcal{G}_i(x_i(t)) \subseteq \Gamma_i \quad (4d)$$

$$\mathcal{G}_i(x_i(t)) \cap \mathcal{G}_j(x_j(t)) = \emptyset, \forall t \geq 0, i, j \neq i \quad (4e)$$

$$\exists T < \infty : \mathcal{G}_i(x_i(T)) \subseteq \mathcal{T}_i, \quad (4f)$$

where $\mathbf{x}(t) = [x_1^T(t), \dots, x_N^T(t)]^T$, $\mathbf{u}(t) = [u_1^T(t), \dots, u_N^T(t)]^T$ represent the states and control signal for each vehicle over the entire operating horizon (i.e., all $t \geq 0$). The OCP is thus the problem of finding the best admissible control inputs $u_i(t)$ for the dynamical systems $f_i(x_i(t), u_i(t), t)$ (4b), starting from the initial conditions $x_{i,0}$, that respect the state constraints (4c), while keeping all vehicles i within their paths Γ_i (4d), avoiding collisions between vehicles (4e), and eventually clearing

the coordination region (4f). The problem captures the basic requirements: safety through (4e), liveness through (4f), and performance through the objective function. Problem 1 can be conveniently reformulated in a discrete time domain by discretizing the systems dynamics (4b). Furthermore, to solve a finite dimensional problem, receding horizon control (RHC) or model predictive control (MPC) schemes can be used [7], where a finite time optimal control problem is solved every sampling time instant. In particular, as explained in Figure 3, RHC seeks future input and state trajectories at every sampling time instant over a finite time horizon, so as to minimize the cost function, subject to the constraints. The first element of the computed control input sequence is applied to the system, and, at the next time step, the problem is formulated and solved over a shifted time horizon. This RHC approach also allows us to account for the future, but by only committing the control action for the current time, we are able cope with limited disturbances (e.g., due to imperfect sensing or communication). This is important, as we will see in the next section.

Challenges in solving the coordination problem

Although finite dimensional, solving problem 1 in a receding horizon framework is extremely challenging, not only from the control perspective, but also due to imperfect communications as well as uncertainties induced by the sensors. While all these challenges are interrelated, we break them down as follows.

Control challenges

The main control-related challenges involve, first, the ability to compute good, feasible control actions, and, second, the ability to guarantee that the closed-loop system has certain desired properties. Regarding the former, note that the mathematical problem of finding the actions of N vehicles that allow them to pass the coordination zone without colliding is inherently a combinatorial problem. For a given initial configuration, a multitude of feasible temporal crossing orders (i.e., different orders in which one vehicle passes a coordination zone before another) might exist, and the optimal ordering can only be found by a structured exploration of the different alternatives. It is therefore no surprise that even the problem of finding a feasible solution to (4) is NP-hard in general [8]. Exact solutions to the OCP are therefore intractable for relevant problem sizes, and either heuristics or approximations must be employed. Regarding the properties of closed-loop control, there are several challenges. For instance, given the severity of constraint violations in the OCP, any controller needs to ensure persistent feasibility. If satisfied, this property ensures that any action taken does not put the system in a state from which no feasible actions exists, i.e., that no vehicle is ever put in a state from which a collision is inevitable. The closed-loop controller must also ensure stability, e.g., to make certain that the crossing order does not change every time the solution is recomputed. Additionally, the aforementioned issues are linked, as the computational challenges of the mathematical

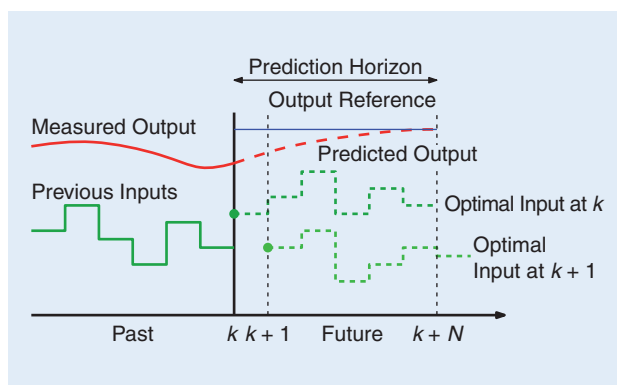


FIGURE 3. An illustration of an RHC scheme. The sketch depicts how, in an RHC algorithm, a present decision is made based on the current state of the system and its predicted future behavior.

coordination problem, for instance, might promote distributed solutions. In that case, the closed-loop controller needs to guarantee the aforementioned properties while the solution is obtained iteratively and possibly asynchronously over the wireless vehicular network.

Communications challenges

Irrespective of how the OCP is solved, information exchange is required between the involved entities (e.g., vehicles and possibly dedicated infrastructure). First and foremost, this includes the information necessary to formulate the OCP, e.g., the models of vehicle dynamics, road geometry, state measurements, and static and dynamic map information. In addition, it also includes information required to solve it, e.g., internal messaging in a distributed, iterative algorithm. Communication between entities will be greatly affected by the impairments associated with wireless channels, including the inherent randomness and correlation of the channel, interference due to simultaneous transmissions, and a limited communication range. In combination with limited communication resources (bandwidth, power), this results in packet drops and random latencies in packet arrivals. For automobile applications, it was pointed out that the current standards for V2V and V2I communication cannot ensure time-critical message dissemination in dense scenarios [9]. In general, it is desired to keep the communication load low, as wireless channel congestion is envisioned to be one of the major challenges related to vehicular networks [10], [11]. Overall, the communication subsystem forms a bottleneck for the OCP, related both to its formulation and the means by which it is solved.

Sensing challenges

The vehicles' own perception of their current locations and the positions of surrounding vehicles is fundamentally uncertain. Both are based on observations from sensors such as cameras, radar, lidar, the global navigation satellite system (GNSS), and inertial navigation sensors, which deliver observations that are corrupted by noise and clutter

(spurious nonobject detection). In addition, the sensors often fail to detect objects, e.g., vehicles and pedestrians, leading to uncertainty about whether all relevant objects are known to the sensing system. Moreover, as each autonomous vehicle is equipped with a different sensor setup with different types of observations, the accuracy of each vehicle's perception of the current traffic situation will typically vary over time and will not be coherent among the vehicles [12]. There are methods to handle time-varying and noncoherent uncertainties, but these require that an uncertainty description is communicated among the involved vehicles, further increasing the demand on the communication system. Even with perfect communication, it is still nontrivial to associate the information from one vehicle with other vehicles' local understanding of the same situation. In the literature, this is called the *data association problem*, which is known to be an NP-hard problem [13]. Processing and sharing a large amount of data requires suitable compression algorithms in combination with application-specific semantic representations, amenable for inclusion in the OCP. These different types of uncertainties are generally ignored in RHC solutions to the OCP, as long as the immediate future is relatively certain. However, in general we may not be able to guarantee performance, liveness, or safety of the resulting solutions. Moreover, the solutions may no longer be stable, which is undesirable from the point of view of the passengers.

Solving the coordination problem

Despite (or perhaps because of) the inherent difficulty of solving the OCP, many solutions have been presented in the context of automated vehicles. These works have been carried out by several communities, resulting in differences of focus and techniques. The resulting techniques can be classified into two groups: rule based and optimization based.

Rule-based solutions

In a large number of existing approaches, e.g., [14]–[16], the vehicle coordination problem is solved using a set of fixed rules, implemented through an interaction protocol. This protocol specifies the content and timing of communications, as well as the possible responses to actions of other participants. To simplify the set of rules, protocols generally assume that individual agents take on partial local responsibility (e.g., resolving rear-end collisions and lane keeping), while a coordination manager resolves any multipath conflicts at the intersection. A canonical protocol operates as follows: 1) a vehicle requests permission to enter the coordination zone at a given time with a given velocity; 2) the intersection manager takes the request and decides whether it can lead to a collision-free crossing (if so, the request is accepted; otherwise, it is denied) 3) when a vehicle's request is denied, it decelerates and sends a new request. Once a request is accepted, the vehicle applies a suitable control

action to meet specifications on when it is allowed to use the coordination zone. From these simple rules, it follows that only requests of vehicles with a safe option are accepted, while all vehicles for which no reservation can be found will slow down and eventually stop.

The benefits of rule-based schemes are the distribution of computation and the economic use of the communication

resources (since the rule-set and interaction protocol is known to all participants). In terms of performance, rule-based solutions are generally only possible to evaluate a posteriori as the actions taken by the vehicles are generated implicitly by application of the rule set. The rules are usually claimed to be chosen to optimize some objective (commonly, throughput), but formal results are missing in most cases.

The general lack of formal guarantees in terms of the objective and constraints of the OCP forms the main weakness of rule-based solutions. Extensions of these works include liveness guarantees [17], and refinements on the individual control policies [18]. The approaches presented in [15], [16], [19], and [20] share similar concepts but differ in terms of the set of rules determining the priority of each vehicle. In summary, while rule-based methods may outperform current regulatory mechanisms, they most likely underutilize the potential of automated vehicles in coordination scenarios.

Optimization-based solutions

In this second group of solution approaches, the coordination problem is treated as a mathematical program from the outset and solved using standard tools and algorithms from optimal control. By doing so, one can potentially separate the feasibility and optimality aspects, and use general, multiobjective performance measures so derive formal guarantees for both performance and safety. However, as a consequence, the computational complexity issues are inherited from the original problem. The contributions of the surveyed papers [21]–[24] are therefore mainly reformulations, approximations, and heuristics that aim to remedy the computational tractability issues. One class of solutions [21], [22] casts an equivalent of the OCP as a safety verification problem. The verification problem entails determining the largest set of (infinite horizon) control actions that avoid any conflict at all future times, and is used in [21] to synthesize a least-restrictive supervisor for human drivers: if the verification fails, the supervisor overrides the human's command, e.g., desired acceleration. Determining the overriding control signals can be posed as a type of OCP wherein the objective corresponds to minimizing the total time needed to clear the intersection or the deviation from the desired control signal, given by the human driver [22]. A more general approach to the OCP was considered in [23], presenting a hierarchical decomposition of (4), where the problem is split up into one centralized time-slot allocation problem and several local vehicle-level optimal control problems. In the latter,

The benefits of rule-based schemes are the distribution of computation and the economic use of the communication resources.

each vehicle computes approximations of its local solutions, parameterized by its occupancy time interval T_i . Loosely speaking, the local optimal cost is expressed as a function of locally feasible T_i 's and transmitted to the central controller. The controller can then find the optimal values of T_i and broadcast this information to the vehicles. Using this approach, the OCP is posed as the search for approximately optimal, nonoverlapping occupancy time slots, which is a rather small mixed-integer optimization problem. Finally, yet another approach was taken in [24], where the combinatorial aspect is resolved through a cooperatively predetermined decision order (or priority), enabling sequential decision making. Once a decision order is agreed upon, the highest priority vehicle solves a local optimal control problem, ignoring all of the remaining vehicles. The solution is communicated to the second vehicle in the ordering, which uses it to solve its own problem: finding the best solution that crosses the coordination zone either before or after the first vehicle. In general, vehicle i in the order will have access to the occupancy intervals of all higher priority vehicles and solves a small local problem. An RHC extension of this approach was developed and demonstrated in [24], where at each time instance both the priority assignment and sequential decision making is repeated.

The primary benefits of the optimization-based approaches are the inherent flexibility and ease with which different and tunable objective functions, dynamics, and physical constraints are included in the design phase, but also modified in the operating phase of the coordination system. This gives to the designer, operator, or passenger the control over what kind of solution the system outputs, and the ability to change this during operation. Furthermore, extensive results regarding the issues of persistent feasibility, stability, and robustness of model-based and optimization-based control schemes (i.e., as in MPC) are available in the literature, as are approximate schemes with quantifiable suboptimality. The optimization-based coordination schemes could potentially leverage such results and formally provide the required safety guarantees. The major weakness of the optimization-based schemes is the complexity, directly inherited from the original formulation (4), which grows exponentially with the number of possible conflict relationships among the vehicles.

The role of signal processing in the OCP

From the aforementioned discussion, it is clear that the OCP explicitly relies on sensing and perception algorithms as well as on wireless communication for its formulation and solution, even though sensing and communication aspects have largely been ignored in the development of control algorithms. Conversely, specific control applications and their demands are usually not considered in the design of sensing

and perception algorithms, nor in the design of wireless communication systems. In this section, we describe recent progress in sensing and wireless communication, and how it relates to solving the OCP. Furthermore, we discuss the need for a tighter integration between the different subsystems, and present ideas on how smart signal processing can be utilized to achieve this.

Wireless communication

Current vehicular communication standards (IEEE Wireless Access in Vehicular Environments and ETSI ITS G5) rely on Wi-Fi-like communication over 10-MHz channels in the 5.9-GHz band and have defined both periodic awareness messages and event-triggered safety messages. These standards can support low-rate (up to 10 Hz) broadcast messages between vehicles within a communication range of about 500 meters [10], but will fail under the high load of the ultrafast communication that is needed to solve the OCP. In contrast, OCP-like problems have been considered explicitly in 5G research [25], with assumed reliability of 99.9% and status updates of 100 ms, considering a steering frequency of 10 Hz. However, these numbers only relate to the dissemination of the final control signal, not the collection of information needed to pose the OCP, nor the iterative message

Signal processing can relieve the burden on communications by censoring less critical information, by tailored compression and semantic algorithms, and by assigning communication resources to those vehicles that are critical to the optimal control problem.

exchange needed to solve it nor do they consider scalability with a large number of vehicles. To get a rough indication for how many vehicles can be supported in a centralized implementation of the OCP, consider a communication system operating in time division multiple access (TDMA) mode, between N vehicles and a controller. This means that each vehicle is assigned a time slot where it during the uplink (UL) phase can transmit its state information to the controller. Assume these time slots last around $100\ \mu\text{s}$, accounting for the actual payload (see Figure 4), as well as overhead (OH) in terms of guard intervals (GIs), training sequences (TSs) and cyclic redundancy check (CRC) bits (for

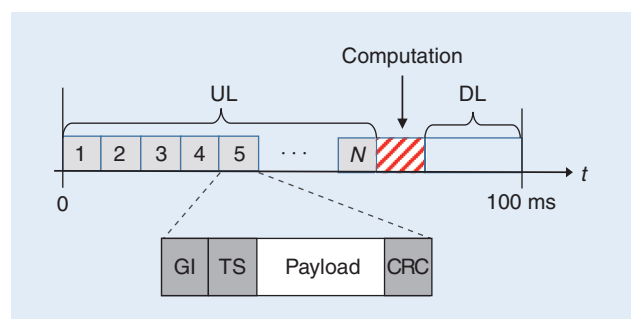


FIGURE 4. An illustration of a TDMA protocol and required communication overhead. In each time step of the RHC, vehicles send state information during the UL phase to the controller, which computes a new control signal and sends this back to the vehicles during the DL phase.

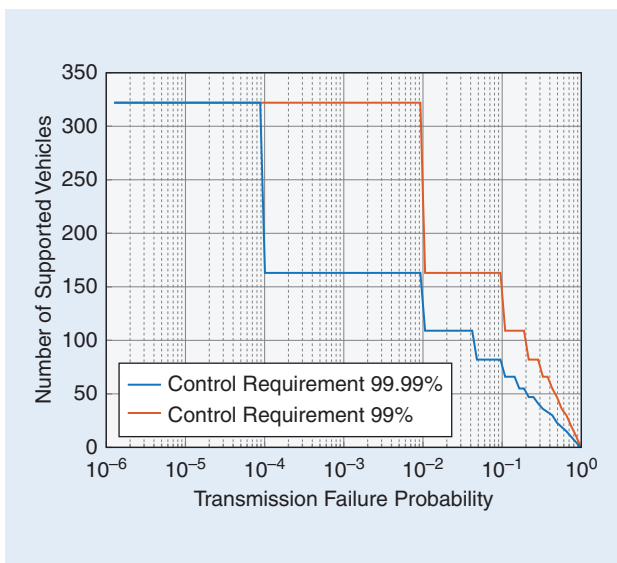


FIGURE 5. The number of vehicles that can be supported in the OCP as a function of the communication failure probability.

comparison, preamble and tail bits in 802.11p adds an overhead of approximately $45 \mu s$). The computation, assumed to scale linearly in N , is set to $10 \mu s$ per vehicle, for some value of the prediction horizon. In the downlink (DL), $200 \mu s$ data

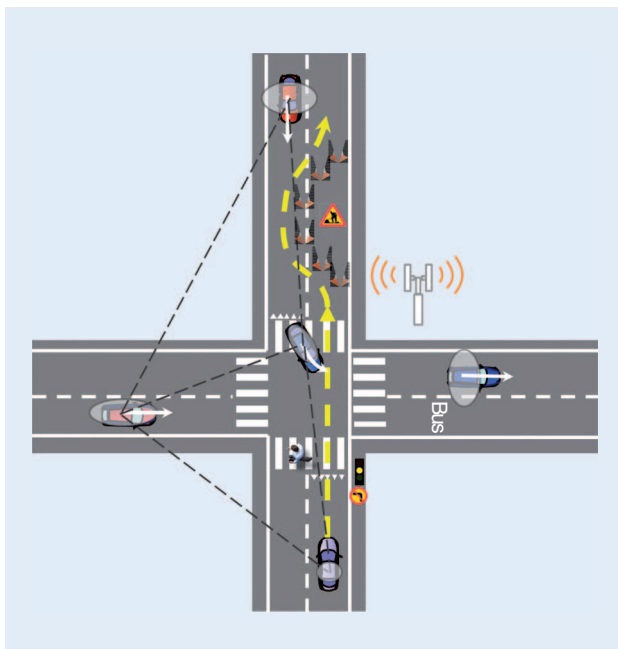


FIGURE 6. An illustration of the sensing problem to support the OCP. The host vehicle (bottom) is approaching an intersection. Information about the intersection is stored in a detailed map containing position of landmarks, geometry of lanes, traffic rules, etc. The aim of the perception subsystem is to position the host vehicle and other relevant road users as well as other obstacles (construction site) in the map such that the OCP can calculate an optimal path (shown as dashed lines). Both the host position and the state of other vehicles (pose and velocities) are described, including uncertainty measures (depicted as gray ellipses). In this example, four of the vehicles are cooperative and exchange information about their positions and current perception to each other as well as the construction site.

packets with the control signals are sent back to each of the vehicles. Since communication can never be guaranteed to succeed, we will consider that the OCP operates under a minimal requirement in terms of the fraction of information that is needed from the vehicles (say, 99% or 99.9%). The communication system is then designed to retransmit data until the requirement is met. Given this information, Figure 5 depicts the number of vehicles that can be supported for different control requirements and different communication failure probabilities. We observe that, when either the channel is very reliable or when the control requirement is loose, then more than 300 vehicles can be supported. However, this number drops quickly when the channel becomes more unreliable. Note that packet error rates in excess of 10% are not uncommon in practice. Based on this quick analysis, it is easy to see that the communication forms a bottleneck for the OCP, especially in urban scenarios where several hundred vehicles can be within communication range. To deal with the conflicting demands on the communication system in terms of latency, throughput, and node density, new V2X communication architectures were proposed in [26]. In addition, dedicated physical layer communication techniques, as well as highly optimized medium access control algorithms, will need to be developed, to complement and support these architectures. Signal processing can further relieve the burden on the communication subsystem by censoring less critical information, by tailored compression and semantic algorithms, and by assigning communication resources to those vehicles that are expected to be the most critical to the OCP. At the same time, signal processing is also expected to play an increasing role in the security of the OCP (through ultrafast authentication and verification), privacy preservation (rendering the OCP and its solutions anonymous and untraceable), and analytics (both within a vehicle and between vehicles, particularly after accidents).

Sensing and perception

To support the OCP, the sensing and perception subsystems have two main goals: 1) to estimate the host vehicle's current location (typically on a highly detailed map) and 2) to determine the position of other road users using noisy sensor observations from onboard sensors such as camera, radar, lidar, and GNSS. Both of these problems are challenging in themselves but can be alleviated by allowing information exchange from cooperating vehicles. Figure 6 depicts an illustration of the problem.

The self-localization problem, in this context, is typically solved by matching current sensor observations of the position of landmarks/features with position of sensor landmarks/features stored in a detailed map. This map is either preconstructed offline and streamed to the vehicle from the cloud, or constructed sequentially and jointly with the estimation of the vehicle's position (simultaneous localization and mapping, also known as SLAM [27]). In the case where the map is pre-constructed, the mapping and the localization problem can be separated, and only the localization

part needs to be solved online [28]. However, offline mapping is time-consuming and may need to be repeated periodically. In contrast, under SLAM, there is no need for offline mapping, rendering it less sensitive to changes in the environment. However, the SLAM problem is inherently more difficult than self-localization in a preconstructed map and thus tends to give inferior positioning accuracy. The problem of estimating the position of relevant road users, including uncertainty measures, is known as a multi-sensor and multiobject tracking problem, which is a well-studied problem within several applications. In contrast to the classical formulation, objects in an automotive setting typically give rise to multiple radar and lidar measurements, thus violating the classical point-source assumption (one measurement per object). Instead, objects such as vehicles, need to be treated as extended objects, which is less studied and typically leads to more complex algorithms. However, including multiple measurements per object also allows for a richer description of the object such as orientation and physical dimensions.

Both self-localization and estimating the position of other road users can be performed cooperatively [29]. For instance, for the latter problem, in addition to exchanging position estimates, information about the physical extension can be shared, thus greatly simplifying the inference and reducing the uncertainty in the position of the objects. However, as the sensor observations are typically not labeled, to use the measurements properly we need to be able to correctly associate them with the information coming from the other vehicles and accurately match them to the local view of the traffic situation, adjusting for delays due to data transmission and asynchronous sensor operation. For self-localization with offline mapping, the map resides in the cloud and can thus easily be shared among the cooperating vehicles. By sharing position estimates in the joint map, together with uncertainty measures, each vehicle can jointly estimate a more accurate ego-position as well as the position of all the other vehicles by fusing with the local perception from the on-board sensors [30]. This way, the self-location problem and positioning of other road users are solved simultaneously. This also leads to the possibility of quickly detecting and sharing changes in the map (e.g., the construction site in Figure 6). To increase the positioning accuracy, estimates of relative position to a selected set of high-quality landmarks can be exchanged between the vehicles and used in a similar manner. For SLAM, cooperation is also beneficial. There are two types of cooperative-SLAM (C-SLAM): centralized and distributed. In the former, the cooperating systems communicate their position estimates and current sensor observations to the cloud where a joint map is formed and shared among the systems [31]. In the distributed versions, however, this information is instead communicated to the individual vehicles, which build and keep their own map using all the information. Both of these C-SLAM methods require that the cooperating entities have a fairly homogeneous sensor setup such that landmarks seen

by one system are also detectable by the other systems. In addition, for the detailed sensors typically used for autonomous vehicles, communicating raw sensor observations is probably not feasible; thus, compression and semantic labeling is needed.

Performance of the OCP in the presence of communication and sensing impairments

To illustrate the role sensing and communication play in solving the coordination problem, we consider an intersection scenario of the type illustrated in Figure 2, where incoming vehicles periodically measure and send their state information (UL) to a centralized controller. The control is performed in a receding horizon fashion, where the controller solves a finite time OCP, and broadcasts the resulting control actions to the vehicles (DL). We simplify the OCP by modeling vehicles as points with positions $x_i(t)$, velocities $\dot{x}_i(t)$, and controls/accelerations $u_i(t) = \ddot{x}_i(t)$ along one-dimensional trajectories, aligned with the center of each road. The intersection is then modeled as an interval $[L_i, H_i]$ on each trajectory. The objective (4a) is chosen to be

$$J_i(x_i(t), u_i(t)) = Q_i \int_0^{t_f} (v_i^{\text{ref}} - \dot{x}_i(t))^2 dt + R_i \int_0^{t_f} u_i^2(t) dt, \quad (5)$$

where v_i^{ref} is a constant reference speed, t_f is a time horizon, and $Q_i > 0, R_i > 0$ are weights set by the user. The liveness constraint (4f) is stated as $x_i(t_f) \geq H_i$ for all vehicles. Finally, the problem is discretized and solved using standard optimization tools.

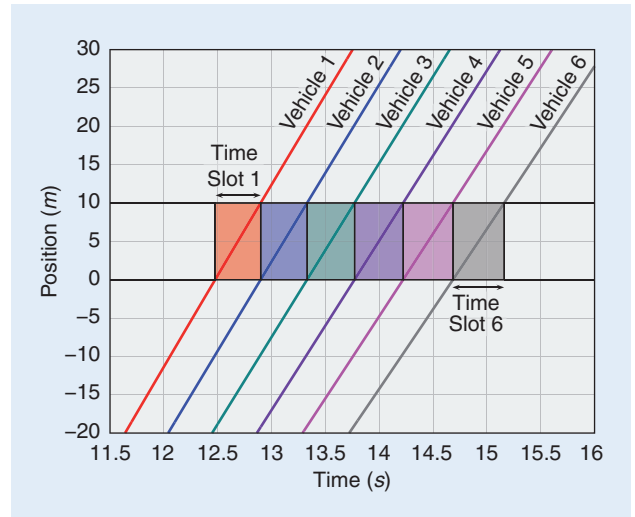


FIGURE 7. The visualization of a part of the position trajectories along each road for six coordinated vehicles under perfect communication and sensing. For each vehicle, the intersection starts at 0 [m] and ends at 10 [m]. The colored lines represent the trajectories of each vehicle. The correspondingly colored boxes visualizes the time slots during which the intersection is occupied by each vehicle. Note that collisions would occur if the time slots were to overlap. In this idealized case, the time slots are tightly packed. Hence, there is no safety margin, and the performance of the system in terms of the objective (5) is pushed to its limits.

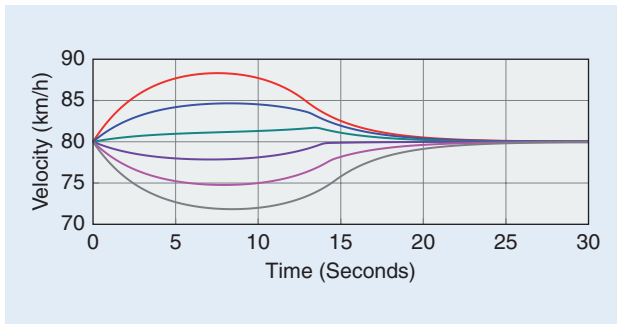


FIGURE 8. The velocity profile for the six coordinated vehicles. Note that the coordination is performed with relatively small adaptations to the velocity of the vehicles well before the intersection is reached by the first vehicle (this happens at around 12.5 seconds, as can be seen in Figure 7).

First, we study an idealized case with perfect communication, no measurement errors, and a perfect match between the dynamics used in the controller and the actual dynamics of the vehicle. We consider an instance of the problem where $N = 6$ vehicles start 300 m away from the intersection at a desired speed of 80 km/hour. Figure 7 shows the solution to the idealized coordination problem in terms of position along their trajectories. We see that the vehicles cross the intersection, one right after the other. Figure 8 shows the velocity profiles of each of the vehicles. The vehicles immediately adjust their speeds to avoid collisions in such a way that minimizes their total cost.

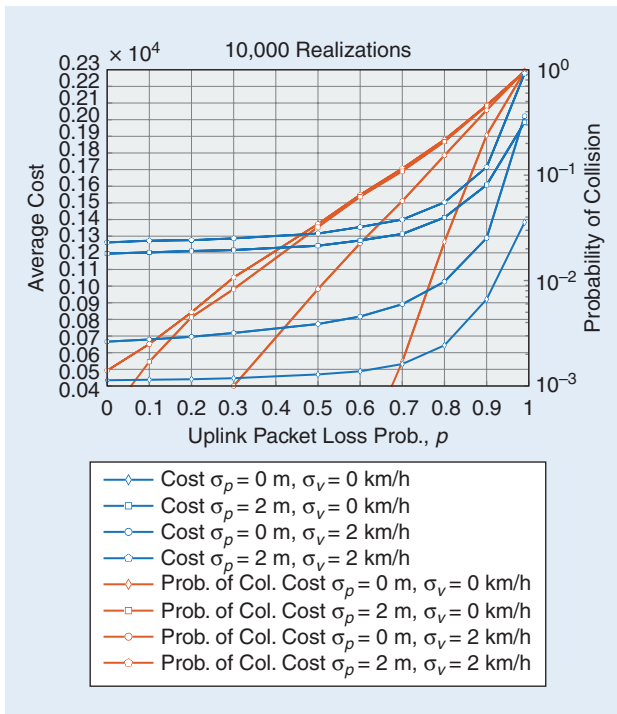


FIGURE 9. The average cost (blue) and probability of collision (red) as a function of UL packet loss probability p for different combinations of sensing uncertainty. From bottom to top: no measurement noise (diamonds), velocity measurement noise (circles), position measurement noise (squares), and both position and velocity measurement noise (pentagons).

To analyze the impact of communication and sensing errors, we reduce the problem size and focus on a two-vehicle case. Both vehicles start 80 m away from the intersection with a speed of 70 km/hour, which will lead to a collision if the central coordinator does not intervene. We introduce a slight mismatch between the true dynamics and the controller model of the dynamics, to avoid degenerate behavior in the presence of packet losses. Packet losses can occur with a probability $p \in [0, 1]$ in the UL communication, while the DL communication is assumed to be perfect. In case a packet is lost, the controller can use the latest received message from a vehicle to predict its current position and speed: the vehicle is simply assumed to obey the previously issued control command. Sensing errors are generated by adding Gaussian noise with a standard deviation σ_p and σ_v to vehicle i 's position and velocity, respectively. Performance is evaluated in terms of 1) the total cost realized by each vehicle and 2) the frequency with which collisions occur. Figure 9 shows the average cost as well as the collision probability, based on 10,000 Monte Carlo runs, as a function of the UL packet loss probability p for different combinations of sensing uncertainty. When there is no sensing uncertainty, we observe that both the probability of collision and the average cost are small, provided p is small (since the true dynamics and OCP model are well matched), but rapidly increase when $p > 0.5$. The reason for this increase is twofold. First of all, it is possible that the coordinator receives the first packet only when the cars are quite close to the intersection, thus requiring more aggressive control and possibly leading to collisions. Second, the controller may operate based on highly outdated information when successive UL transmissions fail, leading to an integration of the mismatch between true dynamics and OCP models, and thus to severe state estimation errors at the controller. When there is sensing uncertainty in either position or velocity, the controller must frequently revise its plan, leading to increases in cost, even under perfect communication. Interestingly, position uncertainty has a higher impact than velocity uncertainty, since velocity uncertainty must be accumulated over multiple failed transmissions to become significant. With increasing packet losses, there is relatively limited impact on the cost, but collisions become more and more frequent. We conclude that even when the OCP has an accurate model for each of the vehicles, sensing uncertainties quickly lead to severe problems, unless the communication system is very reliable. These problems can be avoided by formulating robust versions of the OCP, accounting for worst-case uncertainties, but at a significant cost in terms of the performance.

The road ahead

The coordination problem for cooperative, autonomous vehicles has unique properties compared to other networked control applications due to its safety-critical nature and the challenging communication environment. In this article, we cast such coordination problems as constrained optimal control problems. We have highlighted key challenges in

control, communication, and sensing, along with an overview of recent progress in each of these disciplines. In particular, the theory of distributed optimal control is deemed promising to further develop coordination algorithms that simultaneously accommodate safety and performance. Moreover, uncertainties due to pedestrians and legacy vehicles can be included in such formulations. Nevertheless, scalability and robustness are still challenging problems that deserve further study, in particular, with respect to inherently unreliable exchange of information and limited sensing of the surrounding environment.

Joint design paradigms, where control, sensing, and communications are simultaneously designed, are a promising path forward. While communication-aware control design paradigms exist (i.e., networked control design frameworks), control- and sensing-aware communications need to be further developed to facilitate the solution of the coordination problem. In particular, control-aware communication and sensing systems can establish relationships between the vehicles' mathematical models and the minimal communication and sensing resources necessary to guarantee convergence, stability, and feasibility of the coordination algorithm. Hence, algorithms could be designed so that the coordination plan and the communication or sensing resource allocation are simultaneously decided, thus inherently prioritizing the information exchange of critical vehicles (e.g., vehicles close to the intersection or vehicles that are most likely to be involved in a predicted collision). We believe that joint design of control, communication and sensing systems will pave the way for safer, more efficient, and sustainable road transportation, and that algorithmic aspects associated with signal processing implementations can help address the associated real-world challenges. Proposing and analyzing such joint design frameworks are formidable and long-term research challenges, which will require cooperation among signal processing, communication, and control communities.

Acknowledgments

This work has been supported by the Swedish Research Council under grant 2012-4038; Chalmers' Area of Advance in Transportation; SAFER; the National Metrology Institute hosted at the SP Technical Research Institute of Sweden, which in turn is partly funded by VINNOVA under the program for national metrology (project 2015-06478); COPPLAR (CampusShuttle Cooperative Perception and Planning Platform), funded under Strategic Vehicle Research and Innovation grant 2015-04849; the European Research Council under grant 258418 (COOPNET); the EU Horizon 2020 project HIGHTS (High-Precision Positioning for Cooperative ITS Applications) MG-3.5a-2014-636537; and the grant ADI4VARI02-Progetto ERC BETTER CARS-Sottomisura B.

Authors

Robert Hult (robert.hult@chalmers.se) received his B.S. degree in mechanical engineering in 2011 and M.S. degree in systems, control, and mechatronics in 2013 from Chalmers University, where he is currently pursuing his Ph.D. degree. His current research interests include distributed and cooperative control of automated vehicles and other control topics within intelligent transport systems.

Gabriel R. Campos (gabriel.rodrigues.decampos@polimi.it) received his B.S. degree in electrical engineering in 2007 from the Instituto Superior de Engenharia de Coimbra, Portugal, and from the Université de Reims-Champagne Ardenne, France. In 2009, he received the M.S. degree in electrical engineering from the same university. In 2012, he received his Ph.D. degree in automatic control from the Université de Grenoble, France. He is currently a postdoctoral research associate with the DEIB Department at Politecnico di Milano, Italy. Previously, he was with the Department of Signals

and Systems, Chalmers University of Technology, Sweden. His research interests include cooperative control and autonomous systems.

Erik Steinmetz (estein@chalmers.se) received his M.S. degree in electrical engineering from Chalmers University of Technology, Sweden, in 2009. He is currently a research and development engineer with SP Technical Research Institute of Sweden. He is also affiliated with the Department of Signals and Systems at Chalmers University of Technology, where he is working toward his Ph.D. degree. His research interests include positioning, sensor fusion, communication and controls applied within the fields of intelligent vehicles, cooperative automated driving, and active safety.

Lars Hammarstrand (lars.hammarstrand@chalmers.se) received his M.S. and Ph.D. degrees in electrical engineering from Chalmers University of Technology, Sweden, in 2004 and 2010, respectively. Between 2004 and 2011, he was with the Active Safety and Chassis Department at Volvo Car Corporation, Gothenburg, conducting research on tracking and sensor fusion methods for active safety systems. Currently, he is an assistant professor with the Signal Processing group at Chalmers University of Technology where his main research interests are in the fields of estimation, sensor fusion, and self-localization and mapping, especially with application to self-driving vehicles.

Paolo Falcone (paolo.falcone@chalmers.se) received the Laurea degree in 2003 from the University of Naples "Federico II." He received his Ph.D. degree in information technology in 2007 from the University of Sannio, Benevento, Italy. He is an associate professor with the Department of Signals and Systems of the Chalmers University of Technology, Sweden. He is an associate editor of *Control*

Engineering Practice and *IEEE Transactions on Intelligent Transportation Systems*. His research focuses on constrained optimal control and verification methods, applied to autonomous and semiautonomous mobile systems, cooperative driving, and intelligent vehicles. He is involved in a range of projects, in cooperation with industry, focusing on autonomous driving, cooperative driving, and vehicle dynamics control. He is a Member of the IEEE.

Henk Wymeersch (henkw@chalmers.se) received his M.S. degree in computer science in 2001 and his Ph.D. degree in electrical engineering/applied sciences in 2005, both from Ghent University, Belgium. He is a professor of communication systems with the Department of Signals and Systems, Chalmers University of Technology, Sweden. Prior to joining Chalmers, he was a postdoctoral researcher from 2005 to 2009 with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. He has been an associate editor of *IEEE Communication Letters* (2009–2013), *IEEE Transactions on Wireless Communications* (since 2013), and *IEEE Transactions on Communications* (since 2016). His current research interests include cooperative systems and intelligent transportation.

References

- [1] E.-H. Choi, "Crash factors in intersection-related crashes: An on-scene perspective," *National Traffic Highway Safety Association*, report no. HS-811 366, Washington, D.C., 2010.
- [2] "National motor vehicle crash causation survey," *National Traffic Highway Safety Association*, report no. DOT HS 811 059, 2008.
- [3] M. Campbell, M. Egerstedt, J. P. How, and R. M. Murray, "Autonomous driving in urban environments: Approaches, lessons and challenges," *Philos. Trans. Royal Soc. A*, vol. 368, no. 1928, pp. 4649–4672, 2010.
- [4] S. E. Shladover, "PATH at 20-history and major milestones," *IEEE Trans. Intell. Transport. Syst.*, vol. 8, no. 4, pp. 584–592, 2007.
- [5] R. Di Taranto, S. Muppisetty, R. Raulefs, D. Slock, T. Svensson, and H. Wymeersch, "Location-aware communications for 5G networks: How location information can improve scalability, latency, and robustness of 5G," *IEEE Signal Process. Mag.*, vol. 31, pp. 102–112, Nov. 2014.
- [6] H. Wymeersch, G. R. de Campos, P. Falcone, L. Svensson, and E. G. Ström, "Challenges for cooperative ITS: Improving road safety through the integration of wireless communications, control, and positioning," presented at the *Int. Conf. Computing, Networking and Communications*, 2015.
- [7] E. F. Camacho and C. B. Alba, *Model Predictive Control*. New York: Springer Science, 2013.
- [8] A. Colombo and D. Del Vecchio, "Efficient algorithms for collision avoidance at intersections," in *Proc. 15th ACM Int. Conf. Hybrid Systems: Computation and Control*, 2012, pp. 145–154.
- [9] S. Eichler, "Performance evaluation of the IEEE 802.11p WAVE communication standard," in *Proc. IEEE Vehicular Technology Conf.*, 2007, pp. 2199–2203.
- [10] J. B. Kenney, G. Bansal, and C. E. Rohrs, "LIMERIC: A linear message rate control algorithm for vehicular DSRC systems," in *Proc. 8th ACM Int. Workshop on Vehicular Inter-Networking*, 2011, pp. 21–30.
- [11] E. G. Ström, "On medium access and physical layer standards for cooperative intelligent transport systems in Europe," *Proc. IEEE*, vol. 99, no. 7, pp. 1183–1188, 2011.
- [12] I. Skog and P. Händel, "In-car positioning and navigation technologies survey," *IEEE Trans. Intell. Transport. Syst.*, vol. 10, no. 1, pp. 4–21, 2009.
- [13] J. B. Collins and J. Uhlmann, "Efficient gating in data association with multivariate Gaussian distributed states," *IEEE Trans. Aerospace Electron. Syst.*, vol. 28, no. 3, pp. 909–916, 1992.
- [14] K. Dresner and P. Stone, "Multiagent traffic management: A reservation-based intersection control mechanism," in *Proc. 3rd Int. Joint Conf. Autonomous Agents and Multiagent Systems*, 2004, pp. 530–537.
- [15] H. Kowshik, D. Caveney, and P. Kumar, "Provable systemwide safety in intelligent intersections," *IEEE Trans. Veh. Technol.*, vol. 60, pp. 804–818, Mar. 2011.
- [16] R. Azimi, G. Bhatia, R. Rajkumar, and P. Mudalige, "Ballroom intersection protocol: Synchronous autonomous driving at intersections," in *Proc. IEEE 21st Int. Conf. Embedded and Real-Time Computing Systems and Applications*, 2015, pp. 167–175.
- [17] T.-C. Au, N. Shahidi, and P. Stone, "Enforcing liveness in autonomous traffic management," presented at the *25th Conf. Artificial Intelligence*, San Francisco, CA, August 2011.
- [18] T.-C. Au, M. Quinlan, and P. Stone, "Setpoint scheduling for autonomous vehicle controllers," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2012, pp. 2055–2060.
- [19] K.-D. Kim, "Collision free autonomous ground traffic: A model predictive control approach," in *Proc. ACM/IEEE Int. Conf. Cyber-Physical Systems*, 2013, pp. 51–60.
- [20] M. Ahmane, A. Abbas-Turki, F. Perronnet, J. Wu, A. E. Moudni, J. Buisson, and R. Zeo, "Modeling and controlling an isolated urban intersection based on cooperative vehicles," *Transport. Res. Part C*, vol. 28, no. 0, pp. 44–62, Mar. 2013.
- [21] A. Colombo and D. Del Vecchio, "Least restrictive supervisors for intersection collision avoidance: A scheduling approach," *IEEE Trans. Automat. Contr.*, vol. PP, no. 99, pp. 1–1, 2014.
- [22] G. R. Campos, F. D. Rossa, and A. Colombo, "Optimal and least restrictive supervisory control: safety verification methods for human-driven vehicles at traffic intersections," in *Proc. IEEE Conf. Decision and Control*, 2015, pp. 1707–1712.
- [23] R. Hult, G. R. Campos, P. Falcone, and H. Wymeersch, "An approximate solution to the optimal coordination problem for autonomous vehicles at intersections," presented at the *American Control Conf.*, Chicago, IL, 2015.
- [24] G. R. Campos, P. Falcone, H. Wymeersch, R. Hult, and J. Sjöberg, "A receding horizon control strategy for cooperative conflict resolution at traffic intersections," in *Proc. IEEE Conf. Decision and Control*, 2014, pp. 2932–2937.
- [25] 5G Infrastructure Association, "Automotive vertical sector," White Paper 5G-PPP, 2015.
- [26] H. Laurens, A. Festag, I. Llatser, L. Altomare, F. Visintainer, and A. Kovacs, "Enhancements of V2X communication in support of cooperative autonomous driving," *IEEE Commun. Mag.*, vol. 53, no. 12, pp. 64–70, 2016.
- [27] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: Part I," *IEEE Robot. Automat. Mag.*, vol. 13, no. 2, pp. 99–110, 2006.
- [28] J. Ziegler, P. Bender, M. Schreiber, H. Lategahn, T. Strauss, C. Stiller, T. Dang, U. Franke, N. Appenrodt, C. G. Keller, et al., "Making Bertha drive an autonomous journey on a historic route," *IEEE Intell. Transport. Syst. Mag.*, vol. 6, no. 2, pp. 8–20, 2014.
- [29] S.-W. Kim, B. Qin, Z. J. Chong, X. Shen, W. Liu, M. Ang, E. Frazzoli, and D. Rus, "Multivehicle cooperative driving using cooperative perception: Design and experimental validation," *IEEE Trans. Intell. Transport. Syst.*, vol. 16, pp. 663–680, Apr. 2015.
- [30] A. Howard, M. J. Matarić, and G. S. Sukhatme, "Putting the 'i' in 'team': an ego-centric approach to cooperative localization," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2003, vol. 1, pp. 868–874.
- [31] L. Riazuelo, J. Civera, and J. Montiel, "C2TAM: A cloud framework for cooperative tracking and mapping," *Robot. Auton. Syst.*, vol. 62, no. 4, pp. 401–413, 2014.