



HAL
open science

Safe by Design Autonomous Driving Systems

Marius Bozga, Joseph Sifakis

► **To cite this version:**

| Marius Bozga, Joseph Sifakis. Safe by Design Autonomous Driving Systems. 2024. hal-04829824

HAL Id: hal-04829824

<https://hal.science/hal-04829824v1>

Preprint submitted on 10 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Safe by Design Autonomous Driving Systems

Marius Bozga and Joseph Sifakis

Univ. Grenoble Alpes, CNRS, Grenoble INP, VERIMAG, 38000, France
{mariaus.bozga, joseph.sifakis}@univ-grenoble-alpes.fr

Abstract. Developing safe autonomous driving systems is a major scientific and technical challenge. Existing AI-based end-to-end solutions do not offer the necessary safety guarantees, while traditional systems engineering approaches are defeated by the complexity of the problem. Currently, there is an increasing interest in hybrid design solutions, integrating machine learning components, when necessary, while using model-based components for goal management and planning.

We study a method for building safe by design autonomous driving systems, based on the assumption that the capability to drive boils down to the coordinated execution of a given set of driving operations. The assumption is substantiated by a compositionality result considering that autopilots are dynamic systems receiving a small number of types of vistas as input, each vista defining a free space in its neighborhood. It is shown that safe driving for each type of vista in the corresponding free space, implies safe driving for any possible scenario under some easy-to-check conditions concerning the transition between vistas. The designed autopilot comprises distinct control policies one per type of vista, articulated in two consecutive phases. The first phase consists of carefully managing a potentially risky situation by virtually reducing speed, while the second phase consists of exiting the situation by accelerating.

The autopilots designed use for their predictions simple functions characterizing the acceleration and deceleration capabilities of the vehicles. They cover the main driving operations, including entering a main road, overtaking, crossing intersections protected by traffic lights or signals, and driving on freeways. The results presented reinforce the case for hybrid solutions that incorporate mathematically elegant and robust decision methods that are safe by design.

Keywords: autonomous driving systems · safe by design · assume-guarantee techniques.

1 Introduction

The development of trustworthy autonomous driving systems (ADS) is today a major scientific and technical challenge that could have a considerable economic and societal impact. It will also be an important step towards building systems with human-level intelligence. ADS need to combine extensive cognitive capabilities for situation awareness and management of multiple objectives, while meeting stringent requirements for safety and adaptability.

At present, there is no viable approach allowing to bring satisfactory solutions to this problem. On the one hand, model-based approaches, which we have successfully applied to various safety-critical systems such as flight controllers, are hampered by complexity issues. On the other hand, the AI-based end-to-end solutions already available [29,19] do not offer the necessary guarantees of safety due to the impossibility of explaining neural networks [10]. The immaturity of the state of the art is also confirmed by the current debate on the deployment of autonomous cars in urban environment and real-world experience supporting the view that ADS still have a long way to go before offering acceptable safety guarantees [11].

Currently there is an increasing interest in hybrid design solutions [12,25] combining the best from each approach and integrating ML components, when necessary, e.g., for perception, while using model-based components for goal management and planning. One of the main challenges of hybrid solutions lies in building a model of the vehicle's external environment that accurately describes all the information relevant to its operation. Such a model is essential for deciding which of a set of goals are applicable and, for each of these goals, for generating plans to achieve them. Goal management deals with a variety of possibly conflicting goals subject to different timing constraints. We are distinguishing between short-term goals, such as keeping the vehicle on a trajectory and avoiding collision, long-term goals, such as driving to a destination and medium-term goals concerning the performance of a given driving operation, such as overtaking, merging into a main road or crossing and intersection.

Many studies separate the problem of perception from that of decision-making. They assume that the autopilot has access to a model of the external environment in the form of a map that represents the relevant geometric details of the environment, as well as the signaling equipment and the obstacles around the vehicle, with their positions and kinematic attributes. The formalization of maps and their underlying concepts has been the object of numerous studies including proposals of standards such as [1,2]. Modeling maps as extended graphs is a common idea adopted by many works, one of the main problem being to build maps by composition of heterogeneous elements and data e.g., [20,8,4]. In [6], we study a multilevel semantic framework for the description of maps and their properties, based on metric graphs. The framework allows maps to be specified as a composition of building blocks such as different types of roads and junctions. It also shows how the traffic rules for each of these building blocks induce corresponding safety constraints to be respected by autopilots.

The design of autopilots for autonomous vehicles has been the subject of studies originating in robotics and control theory. Some studies focus on architectures and their hierarchical structure, from the fastest tasks at the lowest level to the slowest tasks at higher levels, e.g. [26]. In the field of architectures, the NIST 4D/RCS reference model [3] provides a basis for the design, engineering and software integration of intelligent systems for unmanned ground vehicles. The reference model integrates planners for a set of tasks, each dedicated to a specific objective for predefined operations. It implicitly assumes that the driving

ability can be summed up as the combination of skills required to perform a set of elementary driving operations. Our work is also based on such an assumption, which allows us to decompose the autopilot behavior into a set of specific control policies that can be designed and verified separately. Validating this assumption poses two problems: the first is to prove the correctness of the composition, i.e., that if the control policies are safe so is the resulting autopilot behavior. The other is completeness, i.e., that the set of control policies considered is sufficient to drive safely in all circumstances.

Of particular interest to this study are correct-by-construction techniques, where the autopilot or some of its features are constructed from a set of properties derived from system requirements. There is considerable work on controller synthesis from a set of system properties, usually expressed in linear temporal logic, see for example [14,22,28,30,31]. These techniques have been extensively studied in the field of control. They consist of restricting the controllable behavior of the system interacting with its environment so that the desired properties are satisfied. However, their practical value is limited by their high computational cost, which depends in particular on the type of properties and the complexity of the system's behavior.

An alternative to synthesis is to achieve correctness by design as a result of composing component properties expressed as "assume-guarantee" contracts specifying a causal relationship between components and their environment: if the environment satisfies the "assume" part of the contract, the state of the component will satisfy the "guarantee" part, e.g., [5,9,17]. A prerequisite for contract-based design is the decomposition of global system requirements into contracts, and the existence of suitable techniques for their implementation [27]. There are a number of theoretical frameworks that apply mainly to continuous or synchronous systems, particularly for analysis and verification purposes [16,21,24]. They suffer from computational limitations because, in the general case, they involve the symbolic solution of fixed-point equations [16]. Furthermore, they only apply to systems with a static architecture, which excludes reconfigurable dynamic systems, such as autonomous systems.

This paper builds on previous results [7] concerning a correct coordination method for autonomous driving systems that allows the construction of a Runtime coordinating a set of autonomous vehicles based on their positions and kinematic attributes. The method considers an ADS as a dynamic system involving vehicles moving in a static environment modelled as a map [6]. A key idea is that the vehicles are bounded to move within their corresponding allocated free spaces computed by the Runtime so as the whole system satisfies given properties including collision avoidance and traffic rules. The main result is that, given a set of traffic rules, the Runtime can calculate authorizations for the coordinated vehicles in such a way that the resulting system behavior respects the traffic rules by design. It is established by showing that the composition of assume-guarantee contracts is an inductive invariant that entails ADS safety.

This work adopts the same perspective of correct-by-design, but considers that there is no overall coordination of vehicles. Instead, they drive au-

tonomously, as in real life, with limited knowledge of their physical environment and surrounding obstacles. The approach adopted is based on the assumption that the complexity of ADS design can be mastered by factoring it into three dimensions.

Locality of context: ADS operate in complex environments taking on a wide variety of configurations, each of which can affect the behavior of the system. Therefore, ADS safety is strongly dependent on the context in which vehicles operate. The traffic infrastructure can be seen as the composition of a finite number of patterns comprising different types of roads and junctions with their signaling equipment. We can therefore imagine that a vehicle’s safety policy is the composition of elementary policies, each of which is used to drive safely according to the corresponding basic road patterns.

Locality of knowledge: A vehicle’s driving policy is based only on local knowledge of the ADS state due to limited visibility. It must therefore drive safely, taking into account the obstacles closest to it, delimited by a visibility zone. In this way, the collective behavior of vehicles in an ADS can be understood and analyzed as the composition of smaller sets of vehicles grouped according to proximity and visibility criteria.

Rights-based responsibility: ADS are a special kind of distributed systems where each agent is responsible for managing a space in its planned route defined by traffic rules. Hence, there is no interaction between the vehicles. Traffic rules guarantee that if each vehicle drives safely in the free space determined dynamically by its rights, then the whole system is safe. This principle of rights-based responsibility[23,13] greatly simplifies the validation problem as the interaction between vehicles is unidirectional (flow-oriented). It is enough to show that each vehicle drives safely in its own free space.

We show that the above decomposition principles reduce the general problem to the design of a vehicle autopilot capable of driving responsibly in a limited number of contexts and configurations involving a relatively small number of other vehicles and objects.

We call *vista* the input of an autopilot generated by its perception function after analysis and interpretation of the information provided by the sensors. Thus, a *vista* defines the state of the environment of a vehicle as well as the applicable traffic rules. The state of the environment includes the positions and kinematic attributes of other vehicles, as well as information on signaling equipment and its status. A *vista* also delimits an area where the vehicle has to act responsibly assuming that the mobile agents in this area behave in accordance with traffic regulations. A vehicle’s *vista* depends on the visibility determined by various factors in its environment, including topology and physical obstacles, as well as weather and light conditions. A simple analysis shows that, as a vehicle moves, its autopilot reacts to inputs that are changes in the state of its environment, characterized by three different types of *vistas*.

1. *road vistas* where there are no crossroads in the vehicle’s area of visibility, and the autopilot is tasked with taking into account the obstacles in its route ahead;

2. *merging vistas* when the vehicle’s route joins a road or a lane where oncoming vehicles have a higher priority and it must therefore give way to these vehicles;
3. *crossing vistas* where the vehicle’s route crosses a junction accessible to other vehicles, and therefore, the vehicle must comply with the traffic rules applicable in this context.

A key idea of this work is that to guarantee safety of an ADS, it is sufficient to ensure safe driving of the vehicles involved for each types of vista types. The main result is obtained by compositionality, in two main steps.

First, we provide safe driving policies for different contexts in each type of vista. Given a reference vehicle, called ego vehicle, safe driving for a vista requires a specific driving operation to overcome potential conflicts while respecting the applicable traffic rules. We argue that critical situations can be characterized by configurations involving, in addition to the ego vehicle, an oncoming vehicle whose route may intersect that of the ego vehicle and a front vehicle located after the intersection on the ego’s route. The vehicle’s driving policy should take into account its dynamic characteristics, in particular its braking and acceleration capacity, as well as the relationships with surrounding obstacles, in particular their speeds and their distances from locations where collisions can occur.

Secondly, based on the assumption that over the course of a journey, the changing perceived reality is a succession of vistas, we show that if a vehicle can drive safely for each vista type, under certain simple temporal conditions, the resulting behavior is safe.

The paper is structured as follows. Section 2 presents the overall approach including definition of basic concepts, the types of vistas and corresponding control policy principles. Section 3 provides control policies for three vista types and demonstrates their safety. Section 4 shows the correctness of the method by proving compositionality of the policies for the different types of vistas. Section 5 concludes by summarizing the main results obtained and discussing future developments with a view to their effective application.

2 The Approach

2.1 Environment modeling with maps

Following the idea presented in [6], we build contiguous road segments from a set \mathcal{S} equipped with a partial concatenation operator $\cdot : \mathcal{S} \times \mathcal{S} \rightarrow \mathcal{S} \cup \{\perp\}$, a length norm $\|\cdot\| : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ and a partial sub-segment extraction operator $[\cdot, \cdot] : \mathcal{S} \times \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathcal{S} \cup \{\perp\}$. Given a segment s , $\|s\|$ represents its length and $s[a, b]$ for $0 \leq a < b \leq \|s\|$, represents the sub-segment starting at length a from its origin and ending at length b . Segments can be used to represent roads at different levels of abstraction. The highest level may ignore the form for the segment and give only its length. The lowest level can be a two-dimensional area. An intermediate level can be a curve showing the form of the road and making abstraction of its width.

We use metric graphs $G \stackrel{\text{def}}{=} (U, E, \mathcal{S})$ to represent maps, where U is a finite set of *vertices*, \mathcal{S} is a set of segments and $E \subseteq U \times \mathcal{S}^* \times U$ is a finite set of *edges* labeled by *non-zero length segments* (denoted \mathcal{S}^*). For an edge $e = (u, s, u') \in E$ we denote $\bullet e \stackrel{\text{def}}{=} u$, $e \bullet \stackrel{\text{def}}{=} u'$, $e.\text{seg} \stackrel{\text{def}}{=} s$. We call a metric graph *connected* if a path exists between any pair of vertices.

We consider the set $\text{Pos}_G \stackrel{\text{def}}{=} U \cup \{(e, a) \mid e \in E, 0 \leq a \leq \|e.\text{seg}\|\}$ of the *positions* defined by a metric graph. Note that positions $(e, 0)$ and $(e, \|e.\text{seg}\|)$ are considered equal respectively to positions $\bullet e$ and $e \bullet$. We denote by $p \xrightarrow{s}_G p'$ the existence of an s -labelled *edge route* between succeeding positions $p = (e, a)$ and $p' = (e, a')$ in the same edge e whenever $0 \leq a < a' \leq \|e.\text{seg}\|$ and $s = e.\text{seg}[a, a']$. Moreover, we denote by $p \overset{s}{\rightsquigarrow}_G p'$ the existence of an s -labelled *route* between arbitrary positions p, p' , that is, $\rightsquigarrow_G \stackrel{\text{def}}{=} (\rightarrow_G)^+$ the transitive closure of edge routes. Finally, we denote by $p' - p$ the *distance* from position p to position p' defined as 0 whenever $p = p'$ or as the minimum length among all segments labeling routes from p to p' or as $+\infty$ if no such route exists. Whenever G is fixed in the context, we will omit the subscript G for positions Pos_G and routes \rightarrow_G or \rightsquigarrow_G .

A connected metric graph $G = (U, E, \mathcal{S})$ can be interpreted as a map, structured into roads and junctions, subject to additional assumptions:

- We restrict to metric graphs which are 2D-consistent [6], meaning intuitively they can be drawn in the 2D-plane such that the geometric properties of the segments are compatible with the topological properties of the graph. In particular, if two distinct paths starting from the same vertex u , meet at another vertex u' , the coordinates of u' calculated from each path are identical. For the sake of simplicity, we further restrict to graphs where distinct vertices are located at distinct points in the plane, and moreover, where no edge is self-crossing (meaning actually that distinct positions (e, a) of the same edge e correspond to distinct points).
- We consider that if the segments of two edges e_1 and e_2 intersect at distances a_1 and a_2 from their starting points, then for the respective positions of the point of intersection we have $(e_1, a_1) = (e_2, a_2)$.

We now define the *junctions* of a map as the classes of an equivalence relation \bowtie on edges, obtained as the transitive closure of the relation generated by pairs of edges (e_1, e_2) such that: either e_1 and e_2 intersect or the endpoint of e_1 and e_2 is the same vertex ($e_1 \bullet = e_2 \bullet$). That is, a junction is a connected sub-graph such that for each edge e_1 there exists another edge e_2 with the same endpoint or intersecting e_1 . Note that if we remove the junctions of a map we obtain sets of roads, where a road is a sequence of non-intersecting edges whose initial vertices have in-degree 1. We assume that junctions are equipped with additional signals to regulate traffic on the edges, e.g., traffic lights, stop signs, etc.

In the remainder of the paper, we consider a fixed metric graph $G = (U, E, \mathcal{S})$ altogether with the junction relationship \bowtie . Also, we extend the junction relationships from edges to their associated positions, that is, consider $(e_1, a_1) \bowtie$

$(e_2, a_2) \stackrel{def}{=} e_1 \bowtie e_2$ whenever $e_1 \bowtie e_2$. Finally, we denote by $r_1 \uplus r_2$ the property that routes r_1, r_2 in G are *non-intersecting*, that is, their sets of positions are disjoint and moreover not belonging to the same junction(s), except for endpoints.

2.2 ADS as dynamic systems

An ADS is a dynamic system involving a set of *vehicles* \mathcal{C} , a set of regulatory *signals* \mathcal{H} , and a map G that represents the environment where the signals are located and the vehicles can move. We use the term *obstacle* to refer to a vehicle or a signal.

The state q of an ADS is the union of the states of its vehicles and signals, $q \stackrel{def}{=} q_{\mathcal{C}} \cup q_{\mathcal{H}}$, where $q_{\mathcal{C}} \stackrel{def}{=} \{q_c\}_{c \in \mathcal{C}}$ and $q_{\mathcal{H}} \stackrel{def}{=} \{q_h\}_{h \in \mathcal{H}}$ where:

- The state of a vehicle c , is a tuple $q_c \stackrel{def}{=} \langle p_c, s_c, v_c, V_c, \dots \rangle$, where p_c is the position of c on the map, s_c is a segment labeling a route on the map starting at p_c , v_c is the speed of c , V_c is its speed limit enforced at position p_c .
- The state of a signal h , is a tuple $q_h \stackrel{def}{=} \langle p_h, type_h, \dots \rangle$, where p_h is its position on the map and $type_h$ denotes its type. We consider signals of the following types:
 - stop or yield signs guarding junctions, with a critical distance attribute cd_h defining the length of the route segment they are protecting,
 - traffic lights guarding junctions, with a critical distance attribute cd_h , a color attribute $color_h$ taking values *red*, *yellow* or *green*, and *time-to-color* attributes $ttred_h, ttgreen_h$; note that $ttred_h = 0$ if $color_h = red$ and otherwise $ttred_h > 0$ is keeping the strictly positive duration until $color_h$ became *red* (same for $ttgreen_h$),
 - speed limits, with an attribute V_h defining the enforced speed limit.

An ADS evolves from an initial state $q[0] \stackrel{def}{=} q_{\mathcal{C}}[0] \cup q_{\mathcal{H}}[0]$ and through states $q[t] \stackrel{def}{=} q_{\mathcal{C}}[t] \cup q_{\mathcal{H}}[t]$, with $q[t] \xrightarrow{\Delta t} q[t + \Delta t]$ where Δt is an adequately chosen time step. The latter can be the period of the autopilots.

In the context of an ADS state, we introduce some additional notations on the underlying metric graphs G . For a vehicle c and non-negative value d we denote by $p_c +_c d$ the unique position p located ahead on the route of c at distance d , that is, formally satisfying $p_c \xrightarrow{s_c[0,d]} p$. Moreover, for any such two positions $p_1 = p_c +_c d_1, p_2 = p_c +_c d_2$ we write $p_1 \sim_c p'_2$ iff $d_1 \sim d_2$ for any $\sim \in \{<, \leq, \geq, >\}$.

Visibility zone: For a reference vehicle called *ego* vehicle, we define the concept of vista characterized by its state and the states of the obstacles in its visibility zone. The visibility zone is defined by an area of the map around the ego vehicle using two types of parameters (Figure 1):




Fig. 1. Vista of an ego vehicle with position p_e and route s_e and two arriving vehicles.

- The *frontal visibility* of the ego vehicle on its route up to a front distance $fd(q_e)$ is delimiting the interval in which the ego's autopilot can perceive the obstacles on its route. The distance $fd(q_e)$ depends on factors such as road curvature, obstacles in view and weather conditions at position p_e .
- The *lateral visibility* of the ego vehicle when its route meets a road or lane r , is the distance from the junction point at which the ego vehicle can perceive vehicles arriving from r . We denote by $ld_r(q_e)$ this distance.

These parameters delimit the visibility zone by points on the route of the ego vehicle and also on the possibly intersecting roads from which vehicles can arrive.

Vistas: Given an ADS, a vista for an ego vehicle with state q_e is a triple $vs \stackrel{def}{=} \langle q_e, q_F, q_A \rangle$ where

- q_e is the state of the ego vehicle.
- q_F is the ordered set of the states of front obstacles $F \stackrel{def}{=} \{f_1, \dots, f_n\} \subseteq \mathcal{C} \cup \mathcal{H}$ located on the route of the ego vehicle in its visibility zone such that $p_e \leq_e p_{f_1} <_e \dots <_e p_{f_n} = p_e +_e fd(q_e)$. We consider that the last visible obstacle f_n is a fictitious vehicle at the front visibility limit.
- q_A is the set of the states of arriving vehicles $A \subseteq \mathcal{C}$, at most one per road joining the route of the ego vehicle in its frontal visibility zone. If there is no real vehicle arriving on the road r within the corresponding visibility distance $ld_r(q_e)$, then we consider a fictitious vehicle a_r with state $q_{a_r} = \langle p_{a_r}, s_{a_r}, v_{a_r}, V_{a_r} \rangle$ defined such that (i) p_{a_r} is the position on the segment representing the road r at distance $ld_r(q_e)$ from the junction point, (ii) s_{a_r} is the segment between p_{a_r} and the junction point and (iii) $v_{a_r} = V_{a_r}$ is the speed limit enforced at p_{a_r} on this route. Thus, for vehicles a_r arriving from a road r , the position p_r of the junction point is such that:

$$\exists d_e \leq fd(q_e). \exists d_r \leq ld_r(q_e). p_e +_e d_e = p_{a_r} +_{a_r} d_r = p_r.$$

Note that while the vista vs includes the states of all visible obstacles in the ego vehicle's route, it only includes the state of a single arriving vehicle

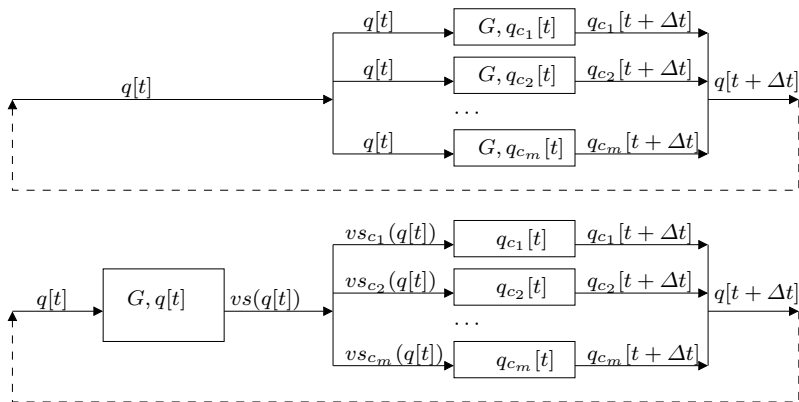


Fig. 2. ADS as a dynamic system composed of the environment and the autopilots.

joining or crossing its route within the frontal visibility limit. The precautionary principle requires us to consider fictitious vehicles at the limits of the visibility zone. These fictitious vehicles can be a frontal obstacle at distance $fd(q_e)$, or a vehicle arriving from a road or lane joining its route at distance $ld_r(q_e)$. In this way, visibility constraints are implicitly taken into account in a vista.

ADS as dynamic systems: An ADS with m vehicles, is a dynamic system using a map G and having states $q[t]$ that can change after time Δt to $q[t + \Delta t]$. At the top of **Figure 2** we show a decomposition of the ADS as a dynamic system with autopilots one for each vehicle c , which knowing the global system state $q[t]$ at time t and the map G , compute their new state $q_c[t + \Delta t]$ at time $t + \Delta t$.

At the bottom of **Figure 2**, this architecture is further refined with the autopilots and a component representing the environment of the vehicles. The environment component receives the states of the vehicles at the end of the cycle and computes the global state of the ADS using the map and the knowledge of the state of its signals. Furthermore, it produces for the c vehicle autopilot the corresponding vista $vs_c(q[t])$ from the global state $q[t]$ taking into account visibility parameters.

2.3 The three basic vista types

As explained in the Introduction, we consider that the autopilot of the ego vehicle receives basic types of vistas as input, each requiring specific operations implemented by the corresponding control policies (**Figure 3**):

① **Road vistas** are of the form $vs \stackrel{def}{=} \langle q_e, q_F, \emptyset \rangle$ where there is no junction in the frontal visibility area of the ego vehicle. Note that road vistas can be simplified in the following manner. Consider the ordered set of the front obstacles F and remind that the last visible obstacle is a fictitious vehicle at the front visibility limit. Clearly, these obstacles are either vehicles or signals (e.g., speed




Fig. 3. Road, merging and crossing vistas - critical sections in gray.

limits). The ego vehicle is responsible for driving safely in the space on its route until to the closest visible vehicle, hence all the obstacles after this vehicle can be omitted as irrelevant.

② **Merging vistas** describe situations where the route of the ego vehicle merges into a main road. They are of the form $vs \stackrel{def}{=} \langle q_e, q_F, q_a \rangle$ such that there exists a merging position p satisfying

$$\exists d_e, d_a, d. p_e +_e d_e = p_a +_a d_a = p \text{ and } s_e[d_e, d_e + d] = s_a[d_a, d_a + d]$$

and moreover $p_e \leq_e p_{f_1} \leq_e p <_e p_{f_2}$ where f_1 is a yield or stop sign. That is, all the elements of F other than f_1 are located after the merging point p . A merging vista can be simplified by replacing q_F by $q_{F'}$ where $\langle q_e, q_{F'}, \emptyset \rangle$ is the simplified road vista for $\langle q_e, q_F, \emptyset \rangle$. To avoid collision, we assume that the signal f_1 is at critical distance cd_{f_1} on the route of the ego vehicle before the merging point.

Note that the ego vehicle deals with merging vistas when its route merges into a higher-priority road or when it changes lanes. In particular, overtaking involves two successive merging operations: one consists of moving from the initial lane to an adjacent lane, the other of returning to the initial lane after a phase of driving in a straight line to ensure that it is far enough away from the overtaken vehicle.

③ **Crossing vistas** describe situations where the route of the ego vehicle crosses a main road. They are of the form $vs \stackrel{def}{=} \langle q_e, q_F, q_A \rangle$ such that there exists a crossing position p satisfying

$$\exists d_e, d_a. p_e +_e d_e = p_a +_a d_a = p$$

for all arriving vehicles $a \in A$. We assume that f_1 is a signal such as a traffic light or a stop sign located before p and protecting a critical distance cd_{f_1} , that is, the length of s_e in the intersection. A simplified crossing vista $svs \stackrel{def}{=} \langle q_e, q_{F'}, q_A \rangle$ is such that $\langle q_e, q_{F'}, \emptyset \rangle$ is a simplified road vista. We can define sub-types of crossing vistas. One is when f_1 is a traffic light, another when f_1 is a yield or a stop sign.

Based on this decomposition into types of vistas, the autopilot can be architected as the serial composition of two components (Figure 4):

- a *Vista Manager* component that receives vistas and produces simplified vistas with their type and relevant parameters such as the protecting signals for crossing vistas and the positions where routes may intersect;

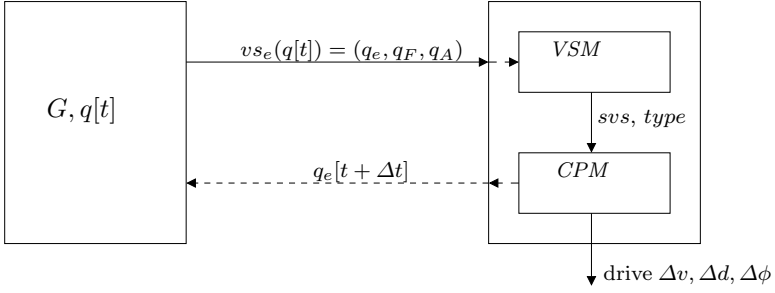


Fig. 4. Architecture of the autopilot.




Fig. 5. A vista and its characteristic parameters.

- a *Control Policy Manager* that receives the simplified vistas and their type, applies the corresponding control policy and produces the next state of the ego vehicle, sent to the environment component responsible for composing the next vista for each vehicle. The Control Policy Manager also generates commands to the drive-by-wire platform to modify the speed by Δv and the steering angle by $\Delta\phi$, in distance Δd . Note that the proposed controllers only calculate the speed variation Δv and distance Δd . We assume that the corresponding steering angle $\Delta\phi$ can be estimated from the curvature of the road segment at the current position and the distance traveled Δd .

2.4 Characteristic parameters of vistas

A vista for the ego vehicle describes a situation involving a potential obstacle on its route and calling for operation implemented by a specific control policy. In its simplest and general form, a vista involves (Figure 5):

- the ego vehicle at position p_e , with its route s_e and speed v_e
- an arriving vehicle a at speed v_a , located at position p_a on its route s_a , which encounters the route s_e at position p
- a signal h at position p_h on the route s_e , located before and protecting the obstacle p with critical distance cd_h i.e., $p_h <_e p \leq_e p_h +_e cd_h$
- a front vehicle f at position p_f on the route s_e located after the obstacle at a distance $d_f = p_f - p$.

It is easy to check that the three different types of vistas match this model. For road vistas the obstacle is just a speed limit signal. Merging vistas cover two cases. The first is when the obstacle can be a yield sign and the road of the ego vehicle merges into a road of higher priority where the arriving vehicle is traveling. The second case is where the ego vehicle moves from one lane to another in which the arriving vehicle is traveling. In both cases, we assume that there is a front vehicle after the merging point. Crossing vistas cover also two cases. The first is when the obstacle is the traffic lights protecting an intersection, in which case arriving vehicles are irrelevant. The second case is an intersection where the crossing ego vehicle faces a yield or stop sign.

3 Control Policies for Vista Types

3.1 Assumptions about the vehicle's dynamics

As explained, an autopilot for a given vista, issues commands that aim at controlling the speed of the vehicle based on its knowledge of the dynamics of the vehicle as an electromechanical system. This knowledge is important for predicting vehicle behavior and its ability to execute given commands aimed at modifying its kinematic state. To avoid detailed modeling of the vehicle as a dynamic system, see for example [15], we assume that we know for each vehicle the next two functions that are sufficient to decide feasibility of control:

1. The *braking function* $B(v)$ that gives the distance needed to brake from speed v to speed 0. We assume the braking function is strict and monotonic, that is, $B(0) = 0$ and $B(v_1) > B(v_2)$ for any $v_1 > v_2$.
2. The *speed control function* $C(v, VL)$ which gives the pair $(\Delta v, \Delta d)$ of the greatest speed variation Δv and the associated distance traveled Δd after Δt , if any achievable from speed v such that $0 \leq v + \Delta v$ and $v + \Delta v, \Delta d$ compliant to speed limit constraints VL as explained below.

It is clear that the speed limit constraints impose changes in speed which must be compatible with vehicle's dynamics characterized by the two functions B and C . We assume that the speed limit constraints VL encountered by a vehicle on its route are specified as sequences of pairs $\langle (d_i, V_i) \rangle_{i=0,n}$ of distance d_i and speed limit V_i (see Figure 6), meaning that the speed limit V_i is enforced in the interval $[d_i, d_{i+1})$ for all $i \in [0, n]$ (and where implicitly $0 = d_0 < d_1 < \dots < d_n < d_{n+1} \stackrel{def}{=} +\infty$). To comply with speed limit constraints, the above-mentioned functions must meet the following requirements:

- First, the C function is only defined for input arguments v and $VL \stackrel{def}{=} \langle (d_i, V_i) \rangle_{i=0,n}$ that satisfy the *controllability* condition $B(v) \leq d_i + B(V_i)$, that is, the vehicle is able to slow down from speed v to V_i in distance less than d_i , for all $i \in [0, n]$.
- Second, in this case, the output pair $(\Delta v, \Delta d) \stackrel{def}{=} C(v, VL)$ shall satisfy, for all $i \in [0, n]$:

- $v + \Delta v \leq V_i$ whenever $d_i \leq \Delta d < d_{i+1}$,
- $\Delta d + \mathbf{B}(v + \Delta v) \leq d_i + \mathbf{B}(V_i)$ whenever $\Delta d \leq d_i$,
- $\Delta d \leq d_i$ whenever $V_i = 0$.

That is, the new speed $v + \Delta v$ satisfies the *controllability* condition with respect to speed limits still ahead after traveling the distance Δd and moreover, Δd never exceeds the distance to a speed 0 limit.

These joint requirements of \mathbf{B} and \mathbf{C} functions allow us later to establish safety invariance properties relating the speed of the vehicle, its braking capacity and its distance to front obstacles (speed limit signals, moving vehicles, etc). Their effective use to define speed control policies for specific driving situations will be explained in the next sections.

As an example, we show how these functions can be defined in a simple case. Assume for the sake of simplicity that the autopilot can only select, during every Δt period, among different constant acceleration/deceleration values in the interval $[-b_{max}, a_{max}] \subseteq \mathbb{R}$, that is, between a maximal deceleration rate $-b_{max}$ and a maximal acceleration rate a_{max} . In this case the braking function \mathbf{B}_{ex} and the speed control function \mathbf{C}_{ex} could be defined as follows:

$$\mathbf{B}_{ex}(v) \stackrel{def}{=} \begin{cases} v\Delta t/2 & \text{if } v < b_{max}\Delta t \\ v\Delta t - b_{max}\Delta t^2/2 + \\ \mathbf{B}_{ex}(v - b_{max}\Delta t) & \text{if } v \geq b_{max}\Delta t \end{cases}$$

$\mathbf{C}_{ex}(v, \langle (d_i, V_i) \rangle_{i=0,n}) \stackrel{def}{=} (a^* \Delta t, v\Delta t + a^* \Delta t^2/2)$ where

$$a^* \stackrel{def}{=} \min_{i=0,n} \max \{ a \in [-b_{max}, a_{max}] \mid v + a\Delta t \geq 0, \Delta d = v\Delta t + a\Delta t^2/2, \\ d_i > 0 \text{ and } \Delta d + \mathbf{B}_{ex}(v + a\Delta t) \leq d_i + \mathbf{B}_{ex}(V_i) \text{ or} \\ d_i = 0 \text{ and } v + a\Delta t \leq V_i \}$$

That is, the function \mathbf{B}_{ex} defines the distance needed to decrease from speed v to speed 0 when only constant decelerations are used for braking during fixed periods Δt . The function \mathbf{C}_{ex} defines the optimal speed variation as the maximal variation satisfying *all* the speed limit constraints.

It is an easy check that, if the inputs given to the function \mathbf{C}_{ex} satisfy the controllability condition $\mathbf{B}_{ex}(v) \leq d_i + \mathbf{B}_{ex}(V_i)$ then, the sets of acceleration/deceleration values over which the maximum is computed are not empty, for every speed limit constraint (d_i, V_i) . That is, the set of values a compliant with the i th constraint contains, at least:

- a_{max} whenever $d_i > 0$ and $v\Delta t + a_{max}\Delta t^2/2 + \mathbf{B}_{ex}(v + a_{max}\Delta t) \leq d_i + \mathbf{B}_{ex}(V_i)$,
- 0 whenever $d_i > 0$ and $v\Delta t + \mathbf{B}_{ex}(v) \leq d_i + \mathbf{B}(V_i)$ or $d_i = 0$ or $v \leq V_i$,
- $-b_{max}$ whenever $b_{max}\Delta t \leq v$,
- $(V_i - v)/\Delta t$ whenever $0 < v - V_i < b_{max}\Delta t$.

For example, **Figure 6** represents the speed limit constraints $VL = \langle (0m, 100km/h), (40m, 50km/h), (140m, 0km/h) \rangle$ with three thick horizontal lines, respectively 100 km/h during 40 meters (interval $[0, 40)$), followed by 50km/h for




Fig. 6. Speed control illustration according to speed limits VL

100 meters (interval $[40, 140)$), and then stop at 0km/h at 140 meters (interval $[140, -)$). For each speed limit, below the solid line of the same color is the corresponding controllability region, assuming $b_{max} = -3.4\text{m/s}^2$ and $\Delta t = 1\text{s}$. Moreover, the area below the dotted line of the same color is the part of the corresponding controllability region from which the vehicle is controllable after applying the maximal acceleration $a_{max} = 2.5\text{m/s}^2$ for Δt .

The figure shows also the possible speed variations and the associated distance for respectively initial speeds of 30, 60 and 90km/h . For initial speed of 30 and 60km/h the vehicle is controllable with respect to the three speed limits, whereas for initial speed 90km/h it is not possible to respect the speed limit of 50km/h . The figure shows in green for 30 and 60km/h , possible choices of speed variations and the corresponding traveled distance during Δt . Note that, whereas every acceleration preserves controllability for initial speed 30km/h , the maximal acceleration $a_{max} = 2.5\text{m/s}^2$ for 60km/h results in a speed exceeding the controllability limit.

3.2 The follow policy and travel time prediction

Using the braking B and speed control C functions, we now define the primitive follow policy and the related travel time prediction function tt . These primitives are used next to define the more sophisticated control policies for vista types.

The $\text{follow}(vs, p)$ policy takes as input a position p located in the current vista vs of the ego vehicle ahead on its route. This policy controls the speed of the ego vehicle so that it follows safely the position p i.e., by keeping the distance $p - p_e$ as small as possible while remaining safe.

First, the policy observes the speed limits enforced on the ego vehicle route in its vista vs till the position p . To do so, it relies on the function $\text{speed-limits}(vs, p)$

defined next to compute the speed limit constraints VL in the appropriate format defined in [subsection 3.1](#).

```

speed-limits( $\langle q_e, q_F, q_A \rangle, p$ )  $\stackrel{def}{=}$ 
  let  $\langle h_1, \dots, h_k \rangle \leftarrow \langle h \in F \mid type_h = \textit{speed limit}, p_h <_e p \rangle$ 
  return  $\langle (0, V_e), (p_{h_1} - p_e, V_{h_1}), \dots, (p_{h_k} - p_e, V_{h_k}), (p - p_e, 0) \rangle$ 

```

That is, the current speed limit V_e is applied at the current position, the next speed limit V_{h_1} is applied from its position p_{h_1} (if any), etc, and the speed limit 0 is applied at position p .

Second, the $\textit{follow}(vs, p)$ policy uses the speed control function C given the current speed v_e and the speed limit constraints VL on the route to compute the commands Δv and Δd for the driving platform (as explained in [subsection 2.3](#) and illustrated in [Figure 4](#)):

```

follow( $vs, p$ )  $\stackrel{def}{=}$  let  $VL \leftarrow \textit{speed-limits}(vs, p)$ 
  drive  $C(v_e, VL)$ 

```

Recalling that the speed control function C is a partial function, note that $\textit{follow}(vs, p)$ executes correctly only if C is called with parameters satisfying the controllability property. To this end, the next proposition establishes the pre/postconditions for the correct execution of $\textit{follow}(vs, p)$ in terms of vs and p according to the assumptions about the B and C functions.

Proposition 1. *Let $\langle h_1, \dots, h_k \rangle$ be the ordered set of speed limit signs from the front obstacles F in vs before p . The policy $\textit{follow}(vs, p)$ executes correctly only if*

1. $v_e \leq V_e$, that is, current speed is lower than the current speed limit,
2. $p_e +_e B(v_e) \leq_e p_{h_i} +_e B(V_{h_i})$, that is, speed limit preservation with respect to any speed limit h_i located such that $p_e <_e p_{h_i} <_e p$,
3. $p_e +_e B(v_e) \leq_e p$, that is, safety distance preservation to the position p .

In this case, the speed variation Δv and the distance traveled Δd satisfy, for any speed limit signal h_i as above (where taking implicitly $p_{h_{k+1}} \stackrel{def}{=} p$):

4. $v_e + \Delta v \leq V_e$ if $p_e +_e \Delta d <_e p_{h_1}$,
5. $v_e + \Delta v \leq V_{h_i}$ if $p_{h_i} \leq_e p_e +_e \Delta d <_e p_{h_{i+1}}$,
6. $p_e +_e \Delta d +_e B(v_e + \Delta v) \leq_e p_{h_i} +_e B(V_{h_i})$ if $p_e +_e \Delta d \leq_e p_{h_i}$,
7. $p_e +_e \Delta d +_e B(v_e + \Delta v) \leq_e p$.

Proof. The $\textit{follow}(vs, p)$ policy calls the speed control function C with the current speed v_e and the speed limit constraints $VL \stackrel{def}{=} \langle (0, V_e), \langle (p_{h_i} - p_e, V_{h_i}) \rangle_{i \in [1, k]}, (p - p_e, 0) \rangle$ to obtain the speed variation Δv . Then, the conditions 1–3 above are simply restating the controllability conditions, that is, the preconditions guaranteeing a valid return value $(\Delta v, \Delta d) \stackrel{def}{=} C(v_e, VL)$, relative to the current speed and the position of the ego vehicles with respect to front obstacles in the vista, respectively:

- $\mathbf{B}(v_e) \leq 0 + \mathbf{B}(V_e)$ (condition 1),
- $\mathbf{B}(v_e) \leq p_{h_i} - p_e + \mathbf{B}(V_{h_i})$ for all $i \in [1, k]$ (condition 2),
- $\mathbf{B}(v_e) \leq p - p_e + \mathbf{B}(0)$ (condition 3)

In a similar manner, the conditions 4–7 are restating the postcondition of the successful call, respectively:

- $v_e + \Delta v \leq V_e$ if $0 \leq \Delta d < p_{h_1} - p_e$ (condition 4),
- $v_e + \Delta v \leq V_{h_i}$ if $p_{h_i} - p_e \leq \Delta d < p_{h_{i+1}} - p_e$ for some $i \in [1, k]$ (condition 5),
- $\Delta d + \mathbf{B}(v_e + \Delta v) \leq p_{h_i} - p_e + \mathbf{B}(V_{h_i})$ for all $i \in [1, k]$ such that $\Delta d \leq p_{h_i} - p_e$ (condition 6),
- $\Delta d + \mathbf{B}(v_e + \Delta v) \leq p - p_e + \mathbf{B}(0)$ (condition 7) because the last limit constraint enforces the speed 0 at distance $p - p_e$. \square

Finally, consider a typical vista vs involving a front vehicle f at position p_f and a signal h located at position p_h , protecting a critical area situated between the ego vehicle and f (as in Figure 5). Usually, the decision for the ego vehicle to cross or not the position p_h between p_e and p_f depends on the time needed to go from p_e to $p_h +_e cd_h$ when safely following f . This travel time can be predicted if we assume the ego vehicle will drive continuously according to the $\text{follow}(vs, p)$ policy at least until it reaches the position $p_h +_e cd_h$.

That is, as $\text{follow}(vs, p)$ is a deterministic policy, we can effectively predict the time to travel from p_e to any position p' such that $p_e <_e p' \leq_e p$ by using the function $\text{tt}(vs, p', p)$ defined below:

```

tt(vs, p', p)  $\stackrel{\text{def}}$  let  $v \leftarrow v_e, VL \leftarrow \text{speed-limits}(vs, p), x \leftarrow p' - p_e, t \leftarrow 0$ 
  while  $x > 0$  do
    let  $(\Delta v, \Delta d) \leftarrow C(v, VL),$ 
       $v \leftarrow v + \Delta v, x \leftarrow x - \Delta d, t \leftarrow t + \Delta t,$ 
       $VL \leftarrow \langle (\max(d_i - \Delta d, 0), V_i) \mid (d_i, V_i) \in VL, \Delta d \leq d_{i+1} \rangle$ 
  return  $t$ 

```

This function “simulates” driving according to $\text{follow}(vs, p)$ and sums up the Δt steps until the ego vehicle travels the distance $x = p' - p_e$. Note that, in this computation we assume that the leading position p is not changing. In typical situations, p may correspond to the position p_f of a front vehicle f moving forward. In any case, however, consequent predictions for the traveling time to p' can either remain unchanged or decrease (but never increase) over time.

3.3 Control policies for vista types

Each vista requires a specific operation when the ego vehicle approaches the obstacle. The aim of the operation is to clear the obstacle safely, respecting traffic regulations and, of course, avoiding accidents with the arriving vehicle and the front vehicle. The operation corresponding to a vista is logically characterized by scenarios comprising two successive phases:

1. A *caution phase* during which the ego vehicle approaches the obstacle, reducing its speed if necessary, and waiting for conditions to be favorable to clear the obstacle, e.g., approaching a merge, approaching a crossing or remaining in the same lane before overtaking.
2. A *progress phase* during which the ego vehicle clears the obstacle after checking that there is no risk of collision with the arriving vehicle or the vehicle in front, e.g., to overtake a vehicle, enter a main road, cross an intersection. The progress phase therefore, consists of moving as quickly as possible to avoid collision with the arriving vehicle, while retaining the possibility of avoiding a collision with the front vehicle.

An important issue in autopilot design is the interplay between cautious behavior and progress. An extremely cautious autopilot may be safe but if it neglects the opportunity to progress it can degrade performance resulting in bad road occupancy and possibly deadlocks. Here are a few examples of over-cautious control policies: driving at low speed on a freeway; stopping before a yield sign even if the priority road is clear; not overtaking a slow vehicle in front, such as a truck, when the ego vehicle's performance allows it and the outside lane is clear; stopping before a green light.

In our design, we make a clear distinction between cautious driving and progress by providing a clearance condition that enables the transition between the two phases. We define control policies as processes run by the autopilot to control the speed of the vehicle by generating every Δt a corresponding speed change Δv . They are defined on the basis of the iterative application of the primitive *follow* policy. Control policies are written as simple programs pol_T , one for every type of vista T , with guarded commands of the form

$$\text{pol}_T(vs) \stackrel{\text{def}}{=} [\text{init} \mapsto \dots] \text{ if } g_1 \mapsto \text{pol}_1(vs, \dots) \mid \dots \mid g_k \mapsto \text{pol}_k(vs, \dots) \text{ fi}$$

Their execution should be understood in relation to the architecture of the autopilot (see [Figure 4](#)) proposed in [subsection 2.3](#). Every Δt , as long as the input vistas vs sent to the autopilot have the same type T , the corresponding policy pol_T executes the command whose guard is true. A guard is a state predicate depending on the states of the vehicles and signals of the input vista. Whenever the type of the input vista vs changes to some $T' \neq T$ at some step Δt , the next policy $\text{pol}_{T'}$ is first (re)initialized and then takes over with the execution of the commands as usual.

3.4 Policy for road vistas

The road policy for a simplified road vista $vs = \langle q_e, q_F, \emptyset \rangle$ is the direct application of $\text{follow}(vs, p_f)$, where f is the last (leading) front obstacle in F i.e., either a real or fictitious vehicle before the frontal visibility limit of the ego vehicle in vs .

$$\text{road}(vs) \stackrel{\text{def}}{=} \text{ if } \text{ true} \mapsto \text{ follow}(vs, p_f) \text{ fi}$$

3.5 Policies for merging vistas

Merging with yield sign The merge-yield policy applies for a simplified merging vista $vs = \langle q_e, q_F, q_a \rangle$ as illustrated in [Figure 7](#).

In the caution phase, the ego vehicle approaches according to $\text{follow}(vs, p_h)$ where h is the yield sign protecting the merging point. That is, the ego vehicle ensures $\mathbf{B}(v_e) \leq d_h$ and can stop if needed at critical distance cd_h before the merging point. In the progress phase, the ego vehicle moves towards the merging point according to $\text{follow}(vs, p_f)$ where f is the leading front vehicle in F .




Fig. 7. Merging vista with yield sign

The clearance condition allowing to switch from the caution to the progress phase is defined as follows. For the arriving vehicle a , we assume that it can drive at the maximal speed limit V_a on its road and it is at distance d_a from the merging point. So, if the ego vehicle decides to progress and so switch to the $\text{follow}(vs, p_f)$ policy, the time to reach the merging point would be at most $\text{tt}(vs, p_h + e\ cd_h, p_f)$. Within this time the arriving vehicle will have traveled distance $V_a \cdot \text{tt}(vs, p_h + e\ cd_h, p_f)$. So the remaining distance from the merging point will be $d_a - V_a \cdot \text{tt}(vs, p_h + e\ cd_h, p_f)$. This distance should be large enough for a safe brake, therefore, $\mathbf{B}(V_a) \leq d_a - V_a \cdot \text{tt}(vs, p_h + e\ cd_h, p_f)$.

The merge-yield policy is therefore specified as follows. Note that the switching between the two phases is controlled by the boolean clearance flag cl initially set to *false* and updated continuously during the caution phase.

```
merge-yield( $vs$ )  $\stackrel{def}{=} \mathbf{init} \mapsto cl \leftarrow false, h \leftarrow \text{yield-sign}(F)$ 
  if  $\neg cl \mapsto \text{follow}(vs, p_h), cl \leftarrow (V_a \cdot \text{tt}(vs, p_h + e\ cd_h, p_f) + \mathbf{B}(V_a) \leq d_a)$ 
  |  $cl \mapsto \text{follow}(vs, p_f)$  fi
```

Lane change The lane change vista $vs = \langle q_e, q_F, q_A \rangle$ is illustrated in [Figure 8](#). In the caution phase, the ego vehicle follows the front vehicle f_1 on its lane. The ego vehicle can initiate the progress phase of the lane change whenever the following conditions hold, respectively:

1. the arriving vehicle a_2 can stop safely behind ego's position, $\mathbf{B}(V_{a_2}) \leq d_{a_2}$,
2. the ego vehicle can stop safely behind the front vehicle f_2 on the other lane, that is, $\mathbf{B}(v_e) \leq d_{f_2}$,

3. neither the front vehicle f_1 nor the back vehicle a_1 are currently performing a lane change.




Fig. 8. Lane change vista

These conditions provide the necessary guarantees for the ego vehicle to (begin to) drive safely in the transition between the two lanes, as required for the lane change. To check condition (3) we assume that vehicles signal lane changes, for example by using flashing lights.

In the progress phase, the ego vehicle follows the front vehicle f_2 while also performing effectively the lane change. This policy is essentially the same as $\text{follow}(vs, p_{f_2})$ but also includes moving safely sideways into the corresponding lane as soon as possible and activating the flashing lights. When the lateral movement is completed, the policy falls back to $\text{follow}(vs, p_{f_2})$. The lane change policy is formalized as follows.

$\text{lane-change}(vs) \stackrel{\text{def}}{=} \mathbf{init} \mapsto cl \leftarrow \mathit{false}$
 if $\neg cl \mapsto \text{follow}(vs, p_{f_1}), cl \leftarrow (\mathbf{B}(V_{a_2}) \leq d_{a_2} \text{ and } \mathbf{B}(v_e) \leq d_{f_2})$
 | $cl \mapsto \text{follow}(vs, p_{f_2})$ **fi**

3.6 Policies for crossing vistas

Crossing with yield sign The crossing vista with yield sign is illustrated in Figure 9. It includes the ego vehicle traveling at speed v_e at a distance d_e from the point of intersection with a main road protected by a yield sign h . In this operation, the ego vehicle should approach cautiously moderating its speed until it decides to progress if there is no risk of collision with some arriving vehicle at distance d_a and with allowed maximal speed V_a . Additionally, it should avoid collision with front vehicles at distance d_f after the intersection. We suppose that around the intersection point there is a critical area delimited by a critical distance cd_h such that the presence of two vehicles in this area is considered a potential accident.

For the caution phase, the ego vehicle is driving according to $\text{follow}(vs, p_h)$, that is, guaranteeing $\mathbf{B}(v_e) \leq d_e$. In the progress phase, the ego vehicle is driving according to $\text{follow}(vs, p_f)$. Hence, the time needed by the ego vehicle to cross and get out of the critical section is $\mathbf{tt}(vs, p_h +_e cd_h, p_f)$. In this time, the arriving vehicle will have traveled maximal distance $V_a \cdot \mathbf{tt}(vs, p_h +_e cd_h, p_f)$. So, the remaining space should be enough to brake if needed to avoid collision in the critical area. This gives the clearance condition $V_a \cdot \mathbf{tt}(vs, p_h +_e cd_h, p_f) + \mathbf{B}(V_a) \leq d_a$




Fig. 9. Crossing vista with yield sign

for switching between the two phases. The crossing with yield policy is therefore formalized as follows:

$$\begin{aligned}
 \text{cross-yield}(vs) &\stackrel{\text{def}}{=} \mathbf{init} \mapsto cl \leftarrow \text{false}, h \leftarrow \text{yield-sign}(F) \\
 &\quad \mathbf{if} \neg cl \mapsto \text{follow}(vs, p_h), cl \leftarrow (V_a \cdot \text{tt}(vs, p_h +_e cd_h, p_f) + \mathbf{B}(V_a) \leq d_a) \\
 &\quad | \quad cl \mapsto \text{follow}(vs, p_f) \quad \mathbf{fi}
 \end{aligned}$$

Traffic lights The crossing vista $vs = \langle q_e, q_F, q_A \rangle$ illustrated in [Figure 10](#) involves the ego vehicle approaching a traffic-light protected intersection. It also involves a signal h of type traffic light with a state variable $color_h$ taking values *red*, *yellow* and *green*.

We assume that we know the duration T_y of the yellow light. Furthermore, we assume that the traffic lights of the intersection have an “*all red*” phase of known duration T_{ar} where all the lights are red before some light passes from red to green. These constants are very important for respecting safety regulations requiring that when the ego vehicle enters the critical section the lights should be either green or yellow. In addition, a vehicle entering the intersection must exit before a light turns green on a transverse road.




Fig. 10. Crossing vista with traffic lights

In the caution phase, the ego vehicle must maintain its speed so that it can stop before the traffic light, i.e. $\mathbf{B}(v_e) \leq d_e$. The clearance condition for the progress phase is twofold:

- First, the ego vehicle should not see red light which means even if the lights switch to yellow right after the decision to cross is taken, the ego vehicle will

reach the entrance of the intersection before the lights turn to red. That is $\text{tt}(vs, p_h, p_f) \leq T_y$.

- Second, the time needed to cross the critical section, that is to reach the position $p_h +_e cd_h$ after the junction must be less than $T_y + T_{ar}$, that is, $\text{tt}(vs, p_h +_e cd_h, p_f) \leq T_y + T_{ar}$.
- Third, the junction is clear of vehicles, that is, no arriving vehicle a is in the junction except if located at its respective entry h_a and not moving $v_a = 0$.

The crossing with traffic lights policy is therefore formalized as follows.

$$\begin{aligned}
 \text{cross-traffic-light}(vs) &\stackrel{\text{def}}{=} \mathbf{init} \mapsto cl \leftarrow \text{false}, h \leftarrow \text{traffic-light}(F) \\
 &\mathbf{if} \quad \neg cl \mapsto \mathbf{follow}(vs, p_h), \quad cl \leftarrow (\text{color}_h = \text{green} \text{ and} \\
 &\quad \text{tt}(vs, p_h, p_f) \leq T_y \text{ and } \text{tt}(vs, p_h +_e cd_h, p_f) \leq T_y + T_{ar} \text{ and} \\
 &\quad \bigwedge_{a \in A} (p_a = p_h \Rightarrow (p_a = p_{h_a} \wedge v_a = 0)) \\
 &| \quad cl \mapsto \mathbf{follow}(vs, p_f) \quad \mathbf{fi}
 \end{aligned}$$

Crossing with all-way stop The crossing vista $vs = \langle q_e, q_F, q_A \rangle$ involves the ego vehicle approaching an intersection protected by a stop sign o as illustrated in [Figure 11](#). We assume the intersection is *all-way stop*, that is, all its entries are actually protected by stop signs.




Fig. 11. Crossing vista with all-way stop

The *cross-stop* policy includes three phases. In the first caution phase, the ego vehicle approaches the junction such that to stop at position of the stop sign p_h . Then, in the second caution phase the ego vehicle waits at the stop position until (i) the intersection is clear of vehicles and (ii) ego is the earliest arrived among the waiting vehicles. Then, in the third progress phase, ego crosses the intersection by following its leading front vehicle f from F .

In order to specify and implement the condition (ii) above, we assume that q_e includes a time variable st_e (*stop time*) that records the exact time when the vehicle stops. Also, we remind that the \bowtie relation holds for positions belonging to the same junction, as introduced in [subsection 2.1](#).

The **cross-stop** policy is therefore formalized as follows:

$$\begin{aligned}
\text{cross-stop}(vs) &\stackrel{\text{def}}{=} \mathbf{init} \mapsto cl \leftarrow \text{false}, h \leftarrow \text{stop-sign}(F) \\
&\quad \mathbf{if} \ p_e <_e p_h \mapsto \text{follow}(vs, p_h) \\
&\quad | \ p_e = p_h \wedge \neg cl \mapsto \text{follow}(vs, p_h), \\
&\quad \quad \quad cl \leftarrow \bigwedge_{a \in A} (p_e \bowtie p_a \implies (p_a = p_{h_a} \wedge v_a = 0 \wedge st_e < st_a)) \\
&\quad | \quad \quad \quad cl \mapsto \text{follow}(vs, p_f) \ \mathbf{fi}
\end{aligned}$$

Note that we tacitly exclude the possibility to observe the same stop time for two vehicles and hence to mutually block the progress over the junction. This can be eventually achieved by using an additional arbitration protocol e.g., a total priority order between junction entries.

4 Correctness of the Method

We will use a compositional approach for guaranteeing safety, regardless the number of vehicles and the map characteristics. For doing that, we first define a concept of *free space* for vehicles depending on their current vista and control policy. Then, we show that the vehicles can adapt their speed depending on their free space, that is, are always driving safely while remaining within their free space. Moreover, we show that the proposed vista management and control policies guarantee the free spaces stay disjoint throughout the execution, and thus ensure the absence of collision between vehicles at any time.

The proof of correctness is, however, subject to few additional mild assumptions, listed here for the sake of completeness:

- (A1) The merging points, crossing points and the junctions of the map are neatly separated by road segments. Thus, the frontal visibility of every vehicle contains at most one critical position i.e., a merging, a crossing or a junction entry-point. Also, no such critical positions exist on the arriving roads, with the exception of the entry points belonging to the same junction, if any. In particular, this assumption implies that the transition from one policy to another is always to/from the **road** policy. Thus, the behavior of the Control Policy Manager can be modeled as a mode automaton where each mode corresponds to the application of one particular type of policy, as illustrated in [Figure 12](#).
- (A2) The frontal and lateral visibility do not retract. That is, $p_e +_e fd(q_e) \leq_e p'_e +_e fd(q'_e)$ and $ld_r(q_e) \leq_e ld_r(q'_e)$ for every vehicle e , at any consecutive positions and states p_e, p'_e, q_e, q'_e respectively, for every arriving road r in its vista. Moreover, a vehicle located at some entry point of a junction has full visibility over all the positions of that junction.
- (A3) Speed limit signals are neatly separated and moreover, the limits never increase when approaching junctions and/or any critical road segments. This assumption is mandatory for the conservative evaluation of the clearance conditions.

- (A4) The time period Δt is small enough to guarantee $v \cdot \Delta t \leq B(v)$, for all $v \geq V_0 > 0$ where V_0 is the minimal speed limit used on roads. For example, if $V_0 = 5km/h$ and $B(v) \sim v^2/(-2b_{max})$ for $b_{max} = -3.4m/s^2$ this assumption reduces to $\Delta t \leq 200ms$. The assumption is mandatory for guaranteeing the vehicles remain in their free spaces at every cycle, as defined next.




Fig. 12. Refined behavior of the Control Policy Manager.

We remind that the behavior of an ADS is characterized by all the possible execution sequences $\langle q[t] \rangle_{t=n\Delta t}$ consisting of successive states reached every Δt . We assume that vehicles operate according to their specific vista policies, as introduced previously. That is, at every step Δt , every vehicle e applies the corresponding caution or progress policy step corresponding to its vista vs .

4.1 Safe vistas and safe states

Let $q \stackrel{def}{=} q[t]$ be an ADS state and e be a vehicle. Let respectively $vs \stackrel{def}{=} vs_e(q) \stackrel{def}{=} \langle q_e, q_F, q_A \rangle$ be the current vista (of some type T) and $pol_T(vs)$ be the control policy applied by e in vs . Consider $q_e \stackrel{def}{=} \langle p_e, s_e, v_e, V_e, \dots \rangle$. Then, we define

- the *front lead obstacle* fo_e as the obstacle from F whose position is followed by e in the current execution step of $pol_T(vs)$, that is, either
 - the signal h protecting the merging or crossing (e.g., yield sign, traffic light, stop sign, etc) during the caution phase, or
 - the last (closest) front vehicle f during the progress phase of $pol_T(vs)$
- the *limit position* $\pi_e \stackrel{def}{=} \min_e(p_{fo_e}, p_e + e \cdot B(V_e), (p_h + e \cdot B(V_h))_{h \in F_{st}})$ where $F_{st} \stackrel{def}{=} \{h \in F \mid type_h = \text{speed-limit}, p_h <_e p_{fo_e}\}$
- the *free space* of e as the route $f s_e \stackrel{def}{=} p_e \overset{s_e[0, \ell]}{\rightsquigarrow} \pi_e$ where $\ell \stackrel{def}{=} \pi_e - p_e$.

In other words, the limit of free space is determined by the strongest constraint resulting either from the position of the front obstacle, or from the distances required to comply with speed limits. Let us observe that $p_e \leq_e \pi_e \leq_e p_{fo_e} \leq_e$

$p_e +_e fd(q_e)$, that is, the free space never exceeds the position of the front lead obstacle, and consequently of the frontal visibility limit of e (because the obstacle fo_e is one of the front obstacles from F , visible in the vista vs , see Figure 13).




Fig. 13. Examples of vistas (gray) and free spaces (yellow)

We call a vista vs of type T *safe* if $I_1(vs) \wedge I_2(vs)$ hold where

$$I_1(vs) \stackrel{\text{def}}{=} (p_e +_e \mathbf{B}(v_e) \leq_e \pi_e)$$

$$I_2(vs) \stackrel{\text{def}}{=} \begin{cases} (fo_e = h \wedge p_e \leq_e p_h) \vee (p_h +_e cd_h <_e p_{fo_e} \wedge \mathcal{C}_T) & \text{if } T \neq \text{road} \\ \text{true} & \text{if } T = \text{road} \end{cases}$$

where \mathcal{C}_T is provided in Table 1, for different types T .

$I_1(vs)$ ensures the speed of the e vehicle is adapted to its free space in its vista vs . $I_2(vs)$ ensures the lead obstacle fo_e is consistently defined with respect to the execution phase of the policy $\text{pol}_T(vs)$ for merging or crossing vista vs . The two terms of the disjunction correspond to the caution and the progress phases respectively. For the caution phase, the safety condition ensures that the free space is limited by the protecting signal h and hence not intersecting the critical section of the junction or merging. In the progress phase, however, the free space may expand over the critical section. In this case, the condition \mathcal{C}_T provides the additional guarantees for having non-intersecting free spaces.

Table 1. Specific safety conditions for merging and crossing vistas

T	\mathcal{C}_T
merge-yield, cross-yield	$V_a \cdot \text{tt}(vs, p_h +_e cd_h, p_{fo_e}) + \mathbf{B}(V_a) \leq d_a$
cross-traffic-light	$p_e \leq_e p_h \implies (\text{tt}(vs, p_h, p_{fo_e}) < \text{ttred}_h) \wedge$ $\bigwedge_{a \in A} \text{tt}(vs, p_h +_e cd_h, p_{fo_e}) < \text{ttgreen}_{h_a} \wedge$ $\bigwedge_{a \in A} p_a \boxtimes p_h \implies (p_a = p_{h_a} \wedge v_a = 0)$
cross-stop	$\bigwedge_{a \in A} p_a \boxtimes p_h \implies (p_a = p_{h_a} \wedge v_a = 0 \wedge st_e < st_a)$

We call a state q *safe* if all the vistas are safe, for all vehicles e .

We call a state q *compliant to speed limits* if $v_e \leq V_e$ for all vehicles e .

Lemma 1. *A safe state is compliant to speed limits and ensures disjoint free spaces for all vehicles.*

Proof. As the vista of every vehicle e is safe it holds $p_e + \mathbf{B}(v_e) \leq_e \pi_e$. By definition of the limit position it holds $\pi_e \leq_e p_e +_e \mathbf{B}(V_e)$. Hence, $\mathbf{B}(v_e) \leq \mathbf{B}(V_e)$ and so $v_e \leq V_e$. Therefore, the state is compliant to speed limits.

By contradiction, suppose two vehicles c_1, c_2 have intersecting free spaces. First, consider that p_{c_2} belongs to the free space of c_1 , that is, $p_{c_1} <_{c_1} p_{c_2} <_{c_1} \pi_{c_1}$. This implies c_2 belongs to the front obstacles of c_1 which is moreover located before fo_{c_1} because $\pi_{c_1} \leq_{c_1} p_{fo_{c_1}}$. But this is impossible by the definition of fo_{c_1} which is either the nearest protecting signal h or the nearest front vehicle to c_1 .

Second, consider there exists a common position p belonging to the two free spaces, that is, $p_{c_1} <_{c_1} p <_{c_1} \pi_{c_1}$ and $p_{c_2} <_{c_2} p <_{c_2} \pi_{c_2}$. This can happen only if p is a merging or crossing position of the routes of c_1 and c_2 . Let T_1, T_2 be the types of vistas of c_1, c_2 respectively. We consider the following cases, the others being symmetric:

- $T_1 = \text{merge-yield}, T_2 = \text{road}$: that is, c_1 must give way to c_2 on the merging at position p . As the free space of c_1 extends beyond p that means c_1 is in the progress phase, thus, following a front vehicle fo_{c_1} located beyond the critical section, that is, after $p_h +_{c_1} cd_h$. Also, c_2 is perceived as an arriving vehicle in the vista of c_1 . Since the merging vista of c_1 is safe it satisfies the associated invariant, hence we have $V_{c_2} \cdot \text{tt}(vs_{c_1}, p_h +_{c_1} cd_h, p_{fo_{c_1}}) + \mathbf{B}(V_{c_2}) \leq d_{c_2}$. This implies $\mathbf{B}(V_{c_2}) \leq d_{c_2}$ and consequently $\pi_{c_2} <_{c_2} p$, thus contradicting the assumption that the two free spaces intersect.
- $T_1 = \text{cross-yield}, T_2 = \text{road}$: identical to the previous case.
- $T_1 = T_2 = \text{cross-traffic-light}$: the free spaces could intersect only if the two vehicles c_1 and c_2 are in the progress phase. If at least one of them is in the junction (i.e., c_1) then the safety condition for the other (i.e. c_2) is violated. That is, the 3rd constraint fails for the vista safety of c_2 whenever $p_{c_1} \bowtie p_h$. If none of c_1 or c_2 has yet entered the junction, their safety conditions would require for both of them to have the traveling time to their respective traffic lights lower than the time to pass to red, that is, the 1st constraint. But this would mean both of the two traffic lights have a non-red color, which is forbidden.
- $T_1 = T_2 = \text{cross-stop}$: the free spaces could intersect only if the two vehicles c_1 and c_2 are in the progress phase. But then, if at least one of them is already in the junction, the safety condition of the other one is violated, as in the previous case. Also, if both are at their entry points, the safety conditions of their vistas would require for both of them to have the lowest stop time, which is impossible. \square

4.2 Preservation of safe vistas

Let respectively $q \stackrel{\text{def}}{=} q[t], q' \stackrel{\text{def}}{=} q[t + \Delta t]$ be two consecutive ADS states. Let e be a vehicle with states $q_e \stackrel{\text{def}}{=} \langle p_e, s_e, v_e, V_e, \dots \rangle, q'_e \stackrel{\text{def}}{=} \langle p'_e, s'_e, v'_e, V'_e, \dots \rangle$ and

vistas $vs \stackrel{\text{def}}{=} vs_e(q) \stackrel{\text{def}}{=} \langle q_e, q_F, q_A \rangle$, $vs' \stackrel{\text{def}}{=} vs_e(q') \stackrel{\text{def}}{=} \langle q'_e, q'_{F'}, q'_{A'} \rangle$ at q and q' , respectively. Let T, T' be the types of the two vistas above and $\text{pol}_T(vs)$, $\text{pol}_{T'}(vs')$ the corresponding control policies.

Let respectively fo_e, fo'_e be the front lead obstacles of e according to its vistas vs, vs' . We call the step from q to q' *non-intrusive for e* if either $p_{fo_e} \leq_e p_{fo'_e}$ or $p'_e +_e \mathbf{B}(V'_e) \leq_e p_{fo'_e}$. That is, either the position of the front obstacle is progressing on the route of e , or it can change arbitrarily as long as its distance to p'_e is greater than $\mathbf{B}(V'_e)$.

We are now ready for proving two key preservation properties. The next two lemmas give the conditions for preservation of vista safety, assuming respectively, the types of the vistas do not or do change between q and q' .

Lemma 2. *If the vista vs is safe, the states q and q' are compliant to speed limits, the step from q to q' is non-intrusive for e and vs' has the same type as vs then, the vista vs' is safe.*

Proof. $I_1(vs')$. The condition $p_e +_e \mathbf{B}(v_e) \leq \pi_e$ is actually equivalent to conditions 1–3 of [Proposition 1](#) when considering the (lead) position $p \stackrel{\text{def}}{=} p_{fo_e}$. Then,

- conditions 4–5 of [Proposition 1](#) guarantee $v'_e \leq V'_e$ depending on the current applicable speed limit
- condition 6 of [Proposition 1](#) guarantees $p'_e +_e \mathbf{B}(v'_e) \leq p_h + \mathbf{B}(V_h)$ for the speed limits $h \in F_{sl}$ such that $p'_e \leq_e p_h <_e p_{fo_e}$
- condition 7 of [Proposition 1](#) guarantees $p'_e +_e \mathbf{B}(v'_e) \leq_e p_{fo_e}$
- the step from q to q' being non-intrusive for e guarantees $p_{fo_e} \leq_e p_{fo'_e}$ or $p'_e + \mathbf{B}(V'_e) \leq_e p_{fo'_e}$.

Using all the above we can infer $I_1(vs')$ that is

$$p'_e +_e \mathbf{B}(v'_e) \leq_e \pi'_e \stackrel{\text{def}}{=} \min_e (p_{fo'_e}, p'_e +_e \mathbf{B}(V'_e), (p_h +_e \mathbf{B}(V_h))_{h \in F'_{sl}})$$

where $F'_{sl} = \{h \in F' \mid h \text{ speed limit, } p_h < p_{fo'_e}\}$. In particular, observe that for any speed limit $h \in F'_{sl}$ that is not taken into account in F_{sl} we have $p'_e +_e \mathbf{B}(v'_e) \leq_e p_{fo_e} <_e p_h <_e p_h + \mathbf{B}(V_h)$.

$I_2(vs')$ First, assume $fo_e = h \wedge p_e \leq_e p_h$, that is, the policy $\text{pol}_T(vs)$ is in the caution phase for vista vs . Two situations can happen during the step:

- *no clearance*: then, at the next vista vs' we still have $fo'_e = h$ and we know from the previous point that $p'_e +_e \mathbf{B}(v'_e) \leq_e p_{fo_e}$. Then, we obtain immediately $p'_e \leq_e p_h$, that is, $I_2(vs')$ holds for the caution phase of $\text{pol}_T(vs')$.
- *clearance*: then, at the next vista vs' we will have $fo'_e = f$ for some front vehicle f located beyond the critical section, that is, $p_h +_e cd_h \leq_e p_f$. Moreover, as the clearance condition holds at vs for $\text{pol}_T(vs)$ we can check that the specific condition \mathcal{C}_T holds at vs' , for every type of merging or crossing vista T . That is, we obtain that $I_2(vs')$ holds for the progress phase of $\text{pol}_T(vs')$.

Second, assume $p_h +_e cd_h \leq_e p_{fo_e} \wedge \mathcal{C}_T$, that is, the policy $\text{pol}_T(vs)$ is in the progress phase at vista vs . Then, we obtain that $p_h +_e cd_h \leq_e p_{fo'_e}$ if $p_{fo_e} \leq_e p_{fo'_e}$. This latter condition holds if the step is non-intrusive and moreover, assuming there are no entry points located immediately after the end of the critical section, that is, Assumption (A1). Also, we can check that the specific condition \mathcal{C}_T holds at vs' as well. In particular, we use that the states q, q' are speed compliant and moreover the speed limits are non-increasing when approaching critical road segments, that is, Assumption (A3). In addition, we use the following monotonicity property of $\text{tt}()$:

$$\text{tt}(vs', p, p_{fo'_e}) \leq \text{tt}(vs, p, p_{fo_e}) - \Delta t$$

for any position p between p_e and p_{fo_e} provided $p_{fo_e} \leq_e p_{fo'_e}$. That is, the time to reach a position can only improve due to the moving forward of the front obstacle. \square

Lemma 3. *If the vista vs is safe, the states q and q' are compliant to speed limits, the step from q to q' is non-intrusive for e and vs' has a different type than vs then, the vista vs' is safe.*

Proof. $I_1(vs')$: as in the previous lemma. $I_2(vs')$: First, consider $T = \text{road}$ and $T' \neq \text{road}$. The change from the **road** policy to any other merging or crossing policy is performed because the first visible obstacle changes from a front vehicle f (in vs) to a signal h (in vs'). This implies $p_f <_e p_h$, otherwise vs would not have been a **road** policy. We know for the **road** policy that $p'_e \leq_e p_f$, that is, the e vehicle must be able to stop behind the front vehicle f . Consequently, $fo'_e = h \wedge p'_e \leq_e p_h$ holds in vs' , that is, $I_2(vs')$ holds on the first case.

Second, consider $T \neq \text{road}$ and $T' \neq T$. That is, the e vehicle is in a merging or crossing vista vs and, while being in the progress phase, it goes beyond the signal h and receives another vista vs' at the next step. By Assumption (A1) we know that all critical regions of the maps (that is, mergings, crossings, junctions, etc) are separated by road segments. Henceforth, $T' = \text{road}$ unless the e vehicle is perceiving another signal h immediately after treating the vista vs , which is impossible. Therefore, $I_2(vs')$ is trivially satisfied as vs' is a **road** vista.

The next lemma gives a property about the evolution of the limit position and the derived free space in a step.

Lemma 4. *If the vista vs is safe, the step from q to q' is non-intrusive for e , then $\pi_e \leq_e \pi'_e$, that is, the free space of e is non-retracting between q and q' .*

Proof. In general, recall that $\min X \leq \min Y$ iff $\forall y \in Y \exists x \in X. x \leq y$. Therefore, let us consider the terms used in the min-definition of π'_e , respectively

- $p_{fo'_e}$: as the step from q to q' is non-intrusive we know that $p_{fo_e} \leq_e p_{fo'_e}$ or $p'_e +_e \mathbf{B}(V'_e) \leq_e p_{fo_{e'}}$. In the first case we are done (as p_{fo_e} occurs in the min definition of π_e). In the second case we are also done as $p_{fo'_e}$ is not the minimal value for the definition of π'_e .

- $p'_e +_e \mathbf{B}(V'_e)$: we have either $V_e = V'_e$, that is, the speed limit does not change, or $V_e \neq V'_e$. In the first case, $p_e +_e \mathbf{B}(V_e)$ occurs in the min definition of π_e and moreover $p_e \leq_e p'_e$. In the second case, $V'_e = V_h$ for some speed limit $h \in F_{sl}$ visible for the ego vehicle e in state q . But then, $p_h +_e \mathbf{B}(V_h)$ occurs in the min definition of π_e and moreover $p_e \leq_e p_h \leq_e p'_e$. Hence, $p_h +_e \mathbf{B}(V_h) \leq_e p'_e +_e \mathbf{B}(V'_e)$.
- $p_h +_e \mathbf{B}(V_h)$ for $h \in F'_{sl}$: we have either $h \in F_{sl}$ or $h \in F'_{sl} \setminus F_{sl}$. In the first case, we are done, as the same term occurs on both sides. In the second case, we have $p_{fo_e} \leq_e p_h <_e p_h + \mathbf{B}(V_h)$. \square

4.3 Preservation of state safety

Let respectively $q \stackrel{\text{def}}{=} q[t]$, $q' \stackrel{\text{def}}{=} q[t + \Delta t]$ be two consecutive ADS states. We say that the step from q to q' has no gaps if for any vehicle e , $p'_e \in fs_e$, that is, the e vehicle moves within its free space.

Lemma 5. *If q is safe then, the step from q to q' has no gaps and is collision-free.*

Proof. Let fix some vehicle e . Remember that the vehicle e is following the front leading obstacle fo_e in the current execution step. Therefore, it must be able to stop at fo_e . We distinguish two cases. First, if the limit position π_e is equal to p_{fo_e} the vehicle must therefore be able to stop at π_e hence, it moves within its free space. Otherwise, the limit position π_e is strictly smaller than p_{fo_e} and we further distinguish two sub-cases:

- $\pi_e = p_e +_e \mathbf{B}(V_e)$: then, the maximal speed of e during the step could be $V_e \geq V_0$, and using Assumption (A4) we obtain $V_e \cdot \Delta t \leq \mathbf{B}(V_e)$ hence, the e vehicle moves in its free space
- $\pi_e = p_h +_e \mathbf{B}(V_h)$ for some speed limit h located ahead on the route of e : then, the maximal speed of e during this step could be $V' \geq V_h \geq V_0$ such that $\mathbf{B}(V') = (p_h +_e \mathbf{B}(V_h)) - p_e$. Again, using assumption (A4) we obtain $V' \cdot \Delta t \leq \mathbf{B}(V')$ which implies that the vehicle e moves within its free space.

In conclusion, the step has no gaps. As q is safe, by Lemma 1 the free spaces are disjoint. Hence, the step is collision-free. \square

Lemma 6. *If q is safe then, the step from q to q' is non-intrusive for all vehicles.*

Proof. By contradiction, assume the step from q to q' is intrusive for some vehicle e , that is, we $p_{fo'_e} <_e p_{fo_e}$ and $p_{fo'_e} <_e p'_e +_e \mathbf{B}(V'_e)$. This means that the front leading obstacle has shifted backwards and is moreover located below the speed limit range $p'_e +_e \mathbf{B}(V'_e)$.

We shall distinguish two cases, depending on the nature of fo'_e . First, consider that fo'_e is a signal h . But then, h was already in the visibility range at state q and, as $p_h < p_{fo_e}$ it has been already considered (i.e., cleared) in the caution phase of the policy.

Second, consider that fo'_e is a vehicle. If fo'_e is the fictitious vehicle located at the front visibility limit then, the visibility limit has retracted from the previous state, where the position p_{fo_e} was visible. This contradicts Assumption (A2). Otherwise, as vehicles are never moving backwards on their routes, fo'_e must be some new vehicle that enters or crosses the road of the e vehicle in the new vista vs' of e at q' . Then, at state q , the vehicle e is approaching a merging, crossing or junction and as $p_{fo'_e} < p_{fo_e}$, it must be in a progress phase by Assumption (A1). We distinguish the following situations depending on the vista type T of e at at state q :

- $T = \text{road}$: then, the vehicle e is driving on a high-priority road, which is crossed or joined by some low-priority road. If some other vehicle fo'_e appears in front of the e vehicle such that $p_{fo'_e} <_e p'_{e'} +_e \mathbf{B}(V'_e)$ that means fo'_e would be also in the progress phase of its vista of type merge-yield or cross-yield. But then, this contradicts the safety of this vista of fo'_e .
- $T = \text{merge-yield}$ or $T = \text{cross-yield}$: this situation is the dual of the above and we have a similar contradiction for the safety of the vista of e , which must not be in the progress phase while another (high-priority arriving) vehicle (that is, fo'_e) could enter on its road.
- $T = \text{cross-traffic-light}$ or $T = \text{cross-stop}$: again, these situations contradict the safety of the vista of e , which means that no other vehicle may cross the junction at the same time.

Lemma 7. *If q is safe then q' is speed compliant and safe.*

Proof. First, by Lemma 1 the state q is speed compliant.

Second, we prove that state q' is also speed compliant. Consider an arbitrary e vehicle and its safe vista vs . As in Lemma 2, the condition $p_e + \mathbf{B}(v_e) \leq \pi_e$ guarantees the conditions 1–3 of Proposition 1 when the e vehicle is following the front leading obstacle fo_e . Then, the conditions 4–5 guarantee $v'_e \leq V'_e$, depending on the applicable speed limit V'_e at the next state q' .

Third, by Lemma 6 the step from q to q' is non-intrusive for all vehicles. We just proven that q and q' are speed compliant. Then, as all vistas vs are safe at q by Lemma 2 and Lemma 3 all vistas vs' are safe at q' . Hence, q' is safe.

Theorem 1. *If the ADS is initially safe, it will remain so throughout its execution. That is, the execution will preserve safety by avoiding collisions and complying with applicable traffic regulations.*

Proof. If the ADS is initially safe, by Lemma 7 all the states along the execution are safe. Moreover, by Lemma 1 these states are compliant to speed limits and collision free. Furthermore, by Lemma 5 the steps between them are collision-free. \square

5 Discussion

This paper makes a fundamental contribution in a rarely explored direction for ADS, where the state of the art in autopilot design focuses on two diametrically

opposed approaches that are unlikely to be accepted in practice. On the one hand, AI-based end-to-end solutions cannot provide the necessary safety guarantees. On the other hand, in addition to inherent complexity limitations, approaches based on controller synthesis over-simplify the problem to make it accessible to mathematical analysis, and are far from being able to realistically take into account the important details linked to various driving operations.

The idea that driving a vehicle boils down to a composition of skills, each dealing with specific situations, is widely held. Advocated by [3], it is also adopted in validation techniques that focus on particular classes of pre-crash scenarios involving high-risk operations [18]. The paper demonstrates the advantages of compositional reasoning, breaking down the general problem into sub-problems that can be tackled by successive realistic simplifications.

A first class of simplifications is based on the application of rules whose scope is limited by context and knowledge, to which is added the principle of rights-based responsibility. This leads to the definition of the notion of vista for each vehicle, and the corresponding free space in which it can circulate in complete safety. The second class of simplifications comes from the fact that analysis of autopilot inputs enables classification into a very limited number of vistas, each characterized by a corresponding control policy. Hence, the idea that the composition according to which safe driving for each type of vista implies safe driving for any route. Finally, a third class of simplifications derives from the assumption that vehicles drive responsibly, strictly observing the traffic rules and staying within their allocated free space in all cases. In this way it is possible to envisage minimal configurations for each vehicle, including a single vehicle in front of it and vehicles whose routes may cross its own. This greatly simplifies the mathematical analysis, which deduces the invariant constraints to be respected by control policies on the basis of knowledge of two functions characterizing a vehicle's controllability. These functions and their properties define a kind of contract between the autopilot and the underlying electrical control system for braking and steering. They are an essential element of the proposed solution, as they provide the predictability on which safe control policies can be built.

The results presented are an important step in a research program into the development of safe *ADS*. They should be complemented by the study of safe policies for other operations such as U-turns and parking maneuvers. In addition, their application needs to be refined by considering more detailed two-dimensional models of the environment and the integration of autopilot with trajectory control modules.

References

1. ASAM OpenDRIVE® - open dynamic road information for vehicle environment. Tech. Rep. V 1.6.0, ASAM e.V. (Mar 2020), <https://www.asam.net/standards/detail/opendrive>
2. ASAM OpenScenario® - dynamic content in driving simulation, UML modeling rules. Tech. Rep. V 1.0.0, ASAM e.V. (Mar 2020), <https://www.asam.net/standards/detail/openscenario>

3. Albus, J.S., et al.: 4D/RCS version 2.0: A reference model architecture for unmanned vehicle systems. Tech. Rep. NISTIR 6910, NIST Publication (August 2022), <https://doi.org/10.6028/NIST.IR.6910>
4. Beetz, J., Borrmann, A.: Benefits and limitations of linked data approaches for road modeling and data exchange. In: EG-ICE. Lecture Notes in Computer Science, vol. 10864, pp. 245–261. Springer (2018)
5. Benveniste, A., Caillaud, B., Nickovic, D., Passerone, R., Raclet, J., Reinkemeier, P., Sangiovanni-Vincentelli, A.L., Damm, W., Henzinger, T.A., Larsen, K.G.: Contracts for system design. Found. Trends Electron. Des. Autom. **12**(2-3), 124–400 (2018)
6. Bozga, M., Sifakis, J.: Specification and validation of autonomous driving systems: A multilevel semantic framework. In: Principles of Systems Design. Lecture Notes in Computer Science, vol. 13660, pp. 85–106. Springer (2022)
7. Bozga, M., Sifakis, J.: Correct by design coordination of autonomous driving systems. Int. J. Softw. Tools Technol. Transf. **25**(5), 625–639 (2023)
8. Butz, M., Heinzemann, C., Herrmann, M., Oehlerking, J., Rittel, M., Schalm, N., Ziegenbein, D.: SOCA: domain analysis for highly automated driving systems. In: ITSC. pp. 1–6. IEEE (2020)
9. Chatterjee, K., Henzinger, T.A.: Assume-guarantee synthesis. In: TACAS. Lecture Notes in Computer Science, vol. 4424, pp. 261–275. Springer (2007)
10. Chen, L., Wu, P., Chitta, K., Jaeger, B., Geiger, A., Li, H.: End-to-end autonomous driving: Challenges and frontiers. CoRR **abs/2306.16927** (2023), <https://doi.org/10.48550/arXiv.2306.16927>
11. Cummings, M.L.: What self-driving cars tell us about AI risks. IEEE Spectrum (July 2023), <https://spectrum.ieee.org/self-driving-cars-266249426>
12. Harel, D., Marron, A., Sifakis, J.: Autonomics: In search of a foundation for next-generation autonomous systems. Proc. Natl. Acad. Sci. USA **117**(30), 17491–17498 (2020)
13. Hasuo, I.: Responsibility-sensitive safety: an introduction with an eye to logical foundations and formalization. CoRR **abs/2206.03418** (2022)
14. Kress-Gazit, H., Pappas, G.J.: Automatically synthesizing a planning and control subsystem for the DARPA urban challenge. In: CASE. pp. 766–771. IEEE (2008)
15. Li, W., Li, H., Xu, K., Huang, Z., Li, K., Du, H.: Estimation of vehicle dynamic parameters based on the two-stage estimation method. Sensors **21**(11) (2021). <https://doi.org/10.3390/s21113711>
16. Mavridou, A., Katis, A., Giannakopoulou, D., Kooi, D., Pressburger, T., Whalen, M.W.: From partial to global assume-guarantee contracts: Compositional realizability analysis in FRET. In: FM. Lecture Notes in Computer Science, vol. 13047, pp. 503–523. Springer (2021)
17. Meyer, B.: Applying "design by contract". Computer **25**(10), 40–51 (1992)
18. Najm, W.G., Smith, J.D., Yanagisawa, M.: Pre-crash scenario typology for crash avoidance research. Tech. rep., National Highway Traffic Safety Administration, US (2007)
19. Nvidia: Nvidia drive overview (2023), <https://resources.nvidia.com/en-us-drive-overview>
20. Poggenhans, F., Pauls, J., Janosovits, J., Orf, S., Naumann, M., Kuhnt, F., Mayr, M.: Lanelet2: A high-definition map framework for the future of automated driving. In: ITSC. pp. 1672–1679. IEEE (2018)
21. Saoud, A., Girard, A., Fribourg, L.: Assume-guarantee contracts for continuous-time systems. Autom. **134**, 109910 (2021)

22. Schwarting, W., Alonso-Mora, J., Rus, D.: Planning and decision-making for autonomous vehicles. *Annual Review of Control, Robotics, and Autonomous Systems* **1**, 187–210 (2018), <https://doi.org/10.1146/annurev-control-060117-105157>
23. Shalev-Shwartz, S., Shammah, S., Shashua, A.: On a formal model of safe and scalable self-driving cars. *CoRR* **abs/1708.06374** (2017)
24. Sharf, M., Besselink, B., Molin, A., Zhao, Q., Johansson, K.H.: Assume/guarantee contracts for dynamical systems: Theory and computational tools. *CoRR* **abs/2012.12657** (2020)
25. Sifakis, J., Harel, D.: Trustworthy autonomous system development. *ACM Trans. Embed. Comput. Syst.* **22**(3), 40:1–40:24 (2023)
26. Staron, M.: *Automotive Software Architectures - An Introduction, Second Edition*. Springer (2021)
27. Sun, M., Bakirtzis, G., Jafarzadeh, H., Fleming, C.: Correct-by-construction: a contract-based semi-automated requirement decomposition process. *CoRR* **abs/1909.02070** (2019)
28. Waqas, M., Murtaza, M.A., Nuzzo, P., Ioannou, P.: Correct-by-construction design of adaptive cruise control with control barrier functions under safety and regulatory constraints (2022), <https://arxiv.org/abs/2203.14110>
29. Waymo: The world’s most experienced driver (2023), <https://waymo.com/>
30. Wongpiromsarn, T., Karaman, S., Frazzoli, E.: Synthesis of provably correct controllers for autonomous vehicles in urban environments. In: *ITSC*. pp. 1168–1173. IEEE (2011)
31. Wongpiromsarn, T., Topcu, U., Murray, R.M.: Receding horizon temporal logic planning. *IEEE Trans. Autom. Control.* **57**(11), 2817–2830 (2012)