

Modeling Concurrency and Reconfiguration in Vehicular Systems: A π -calculus Approach

Joseph Campbell, Cumhuri Erkan Tuncali, Peng Liu, Theodore P. Pavlic, Umit Ozguner and Georgios Fainekos

Abstract—As autonomous or semi-autonomous vehicles are deployed on the roads, they will have to eventually start communicating with each other in order to achieve increased efficiency and safety. Current approaches in the control of collaborative vehicles primarily consider homogeneous simplified vehicle dynamics and usually ignore any communication issues. This raises an important question of how systems without the aforementioned limiting assumptions can be modeled, analyzed and certified for safe operation by both industry and governmental agencies. In this work, we propose a modeling framework where communication and system reconfiguration is modeled through π -calculus expressions while the closed-loop control systems are modeled through hybrid automata. We demonstrate how the framework can be utilized for modeling and simulation of platooning behaviors of heterogeneous vehicles.

I. INTRODUCTION

The DARPA Grand Challenges and, in particular, the Urban Challenge, demonstrated the feasibility of fully autonomous vehicles driving in urban and rural areas. Since then, multiple well-established companies, research labs, and startups are competing towards becoming the first to sell fully autonomous vehicles capable of driving on the same roads as human-operated vehicles. One important question that has not yet been addressed with current research and development activities is how to enable collaboration and cooperation among the autonomous (and even semi-autonomous) vehicles.

Cooperation is essential in many practical applications of autonomous vehicles, such as platooning. A vehicle platoon is a formation in which several vehicles closely follow each other in order to reduce aerodynamic drag, yielding both reduced fuel consumption [1] and increased road capacity [2]. Another typical application where vehicle cooperation is required is collision avoidance at intersections [3]. In this application, when two or more vehicles approach an intersection from different directions, they communicate directly or indirectly in order to cross the intersection without (usually) coming to a full stop.

Prior work has tackled the challenge of vehicle cooperation by focusing on the development of control and scheduling algorithms for homogeneous vehicles with simplified

dynamics. Another common assumption in the current literature is to ignore any potential issues in the communication protocols of the vehicular systems. Both assumptions can be very limiting when trying to utilize such control algorithms on real networks of heterogeneous vehicles operating at high speeds.

In this paper, we develop a modeling framework for collaborating vehicular systems where both the high-level communication protocols and the low-level complex vehicle dynamics can be modeled in the same framework. For modeling the communication and reconfiguration layer of the system we have chosen the formalism of π -calculus, which was specifically developed for modeling and reasoning over mobile communicating processes where the network structure can be modified dynamically [4]. For the closed-loop system dynamics of each vehicle, we use the modeling framework of hybrid automata [5]. The two layers communicate and synchronize through message passing.

The primary and immediate benefit of such a framework is its flexibility. Different communication protocols can be modeled quickly in a hierarchical fashion where at the highest level we can potentially verify correctness of the protocols. For example, the proposed framework can easily incorporate verifiable privacy protocols [6] or broadcasting protocols [7]. Furthermore, different vehicles and control algorithms can be easily modeled and simulated because hybrid automata model both the control and continuous dynamics of the vehicles as well as any other discrete modes of operation for these vehicles (e.g., emergency braking, economy versus sport driving mode). The main benefit of utilizing hybrid automata as a modeling framework is that reachability analysis [8] and automated test generation [9] can be performed efficiently and effectively. At the same time, executable code can be automatically generated for the lower levels [10].

Because autonomous vehicles are safety-critical systems, there needs to be a requirements-based analysis of the whole system rather than of specific components (e.g., [7]) or behaviors (e.g., [11]). Ultimately, there should be a verification framework for heterogeneous, high-fidelity models of collaborating vehicular systems similar to what was proposed by Damm et al. [12] for the Automatic Train Protection system. In particular, stable and safe vehicle controllers for *simplified* models can be tested for safety violations using *high-fidelity* vehicle models in automated test generation tools such as S-TALIRO [9]. This work is the first step

J. Campbell, C.E. Tuncali, T.P. Pavlic, and G. Fainekos are with SCIDSE, Arizona State University, Tempe, AZ 85281, USA. {jacampb1, etuncali, tpavlic, fainekos}@asu.edu

P. Liu and U. Ozguner are with ECE, Ohio State University, Columbus, OH 43212, USA. {liu.3193, ozguner.1}@osu.edu

This work has been partially supported by award NSF CPS 1446730 and NSF 1446735.

towards that goal.

In summary, our contribution in this paper is the development of a modeling framework that can represent complex vehicular networks. The complexity enters both in terms of vehicle dynamics and in terms of complex communication and decision-making protocols. We show that our framework can effectively model heterogeneous vehicular networks running a variety of control algorithms. Finally, the proposed framework lays the groundwork in making verification (at least in terms of automatic test generation) of these systems both tractable and practical.

II. RELATED WORK

A great deal of effort has gone into proving the safety of autonomous vehicle systems. Due to the large and diverse literature, we will present a few works that have inspired our own approach. The reader can find further references in the literature discussed below.

Some methods, such as those by Asplund et al. [13], use formal methods to verify cooperation among autonomous vehicles. In particular, the authors formalize a distributed coordination protocol using finite state machines, and they consider an over-approximation of the vehicles' dynamics based on the maximum acceleration of the vehicles. Then, they utilize Satisfiability Modulo Theory (SMT) solvers to prove safety of an intersection collision avoidance protocol. The results cannot be generalized since they apply only to scenarios where the intersection is a shared resource where only one vehicle has access at each time.

Lygeros et al. [14] model platooning as a hybrid automaton, and they automatically design control laws for provably safe merge and split vehicle operations. However, their work does not look into the formalization of the coordination layer for the vehicles and instead treats it rather informally. Our work can be thought of as enabling the modeling of more complicated communication and control algorithms on top of or as an extension of the hybrid automaton models proposed by Lygeros et al. Beyond more complex coordination protocols that include privacy and security, we can study for example the behavior of the platoons when there is an unexpected obstacle on the road.

More recently, Loos et al. [15] apply theorem proving methods to verify the safety of adaptive cruise controllers. In particular, they model the case of vehicles with known bounds on important characteristics (i.e., maximum braking deceleration, maximum acceleration, and worst-case response-time) with cruise-control implementations which pick an acceleration subject to a special safety condition. Under that condition, they show that the resulting arbitrarily large group of vehicles will be collision free. A similar result is shown by Platzer [16] for the case of vehicles entering a highway from an on-ramp, where collision-free safety can be proved under the assumption that vehicles entering the highway meet a condition based on their position, speed, and acceleration relative to vehicles on the highway. Although it is hinted that this framework can capture the effect of a communication-mediated reconfiguration of the vehicle

platoon, neither case explicitly models any communication nor coordination protocols between vehicles. Platzer [16] clarifies that communication can be modeled within their framework, but shared variables with access delays must be used as opposed to explicitly modeling communication channels that are opened, maintained, and closed by vehicles. Furthermore, despite the expressiveness of the framework, direct and automated synthesis of controllers that meet required specifications remains elusive.

Along another direction to this problem, Franzle et al. [17] present an extension to the Multi-Lane Spatial Logic (MLSL) [18] by introducing a local scope for the observations of each vehicle. MLSL has been proposed as a logical framework for reasoning about decision-making algorithms for automated driving.

A number of works consider the verification and synthesis problem for cooperative vehicle behaviors exclusively at the supervisory or communication level only. That is, no continuous vehicle dynamics are explicitly considered. For example, Bochmann et al. [19] present a discrete-event controller synthesis algorithm for lane-changing maneuvers. Bengtsson et al. [20] present a communication protocol for enabling two platoons to merge together. The work of Kamali et al. [21] is very similar to our proposed architecture in the sense that they also enable simulation of collaborative vehicles by utilizing high-fidelity models at the physical level. However, the vehicle coordination protocol is modeled through timed automata instead of mobile process calculi.

Finally, we remark that some preliminary results of our work were presented in a news brief [22]. The news brief primarily argued on the benefits of utilizing π -calculus in the modeling of collaborative vehicular systems. Here, we extend upon that announcement by presenting details on the π -calculus formalism, the low-level control layer, the interaction of high- and low-level layers, and the experimental results.

III. PRELIMINARIES

In this paper, we model concurrent high-level behavior of autonomous vehicles with a process algebra known as π -calculus. This section provides a brief overview of the algebra along with an extension that is necessary to fully represent a distributed autonomous system.

A. Process Algebra

Process algebras, or process calculi as they are also known [23], are a way to algebraically model the interaction between concurrent systems. They are exceedingly useful for reasoning about parallel systems and as such allow for the verification of these systems to determine whether certain properties hold. For example, if we are considering a system of autonomous vehicles that are traveling together in a platoon, then we would like to reason about such a system and establish that merging behavior protocols never produce a deadlock.

To this end, we must first establish the notion of a process. Informally speaking, a process is simply the behavior of

a system [23]. Process algebra puts forward the idea that these behaviors can be modeled as a sequence of actions over time which can be manipulated with algebraic rules. As a simple example, consider an abstracted model of a process for a vehicle that is changing lanes. Given the actions, *enable_signal*, *disable_signal*, *change_lane*, and the sequential operator (\cdot) this process can be modeled as *enable_signal.change_lane.disable_signal*. This expression means that the observed behavior is to enable the turn signal, change lanes, and finally disable the turn signal. However, as we describe in more detail in the following section, the real power of process algebra comes from its ability to describe the interaction between concurrent systems.

B. π -Calculus

In this work, we employ a specific variant of process algebra known as π -calculus [24], [4]. In its most basic form, π -calculus defines an expression for a process $P \in \mathcal{P}$ as a sequence atomic actions π . The syntax of a process expression can be described by the following grammar.

$$P ::= \pi.P \mid 0 \quad (1a)$$

Semantically, this means that the process P is empty (0) and thus terminating or is composed of one or more atomic actions π that execute in sequence and yield another process expression. Atomic actions in π -calculus are made up of primitives known as *names*. As defined in [24], these may be variables, unobservable actions, communication channels, or simply data values; there is no distinction between them. In this work, we make use of communication channels. In particular, given an infinite set of names \mathcal{N} and $w, x, y, z \in \mathcal{N}$,

- Outgoing communication occurs when a name is sent along out over a channel. The expression $P = \bar{x}\langle y \rangle$ defines P as sending the name y out over x . In this example, x represents a communication channel to another process, and P is sending the message y over that channel.
- Incoming communication occurs when a name is received along a channel. The expression $P = x(z)$ defines P as receiving the name z from x . Similar to the previous example, x represents a communication channel to another process, and P receives a message on that channel which gets bound to the name z .

In addition, process expressions may take a parametric form. For example, $P(x) = \bar{x}\langle y \rangle$ is a parametric definition of process P . When P is referenced, it must be passed a name. If we have the expression $w(z).P(z)$, a message is received over channel w that gets bound to z and is then passed to P . In this context, z must be another communication channel itself because P immediately uses it to send message y . This ability to transmit and receive communication channels is an important concept in π -calculus and is referred to as *mobility*. Furthermore, π -calculus may be extended to accommodate multiple parameters; this is known as polyadic π -calculus [4].

We now introduce more complex constructs into the syntax, as summarized by the following grammar.

$$P ::= P \mid P \parallel P' \mid P + P' \mid !P' \quad (1b)$$

$$P ::= P \mid (\nu x)P \mid x : [y \Rightarrow P, z \Rightarrow P'] \quad (1c)$$

These constructs support the modeling of multiple processes. In particular,

- $P = P \parallel P'$ defines the process P as concurrently executing P and P' .
- $P = P + P'$ defines the process P as non-deterministically executing either P or P' .
- $P = !P'$ indicates that there are as many concurrently executing copies of P' as desired [4].
- $(\nu x)P$ indicates the name x is unique to P .
- $P = x : [y \Rightarrow P, z \Rightarrow P']$ defines the process P as resulting in P if the name x is equal to y , or resulting in P' if x is equal to z . Names can also be constants, and so it is possible define $P = x : [\text{True} \Rightarrow P, \text{False} \Rightarrow P']$ for evaluating against True and False.

Now that the syntax and an informal explanation of the semantics for concurrent processes have been introduced, the concept of *reaction* [4] must be discussed. Consider two processes, $P = x(y).P'$ and $Q = \bar{x}\langle z \rangle.Q'$, and suppose they are running concurrently, $P \parallel Q$. The names x and \bar{x} are considered complementary, and the actions are linked; if z is sent over x by Q , then it will be received by P and bound to y .

This highlights the importance of the scope restriction operator, ν . Drawing on an example from Milner [25], let $P(x) = \bar{x}\langle y \rangle$, $Q(x) = x(z)$, and $R(x) = x(w)$. If the processes run concurrently as in $P(x) \parallel Q(x) \parallel R(x)$, then there are two outcomes: either Q receives the message sent by P , or R does. This is a consequence of all three processes sharing the same x channel. However, if x is now restricted as in $(\nu x)(P(x) \parallel Q(x)) \parallel R(x)$ then there is only one outcome: Q receives the message sent by P along their shared (and unique) channel x .

C. ω -Calculus

This work also makes use of ω -calculus [7], an extension to π -calculus that introduced semantics for reasoning about mobile ad hoc wireless networks. It is a particularly useful extension when dealing with distributed robotic systems that must form ad hoc networks in order to establish communication. We draw from two operations that have been previously introduced in ω -calculus:

- $\bar{\mathbf{b}}\langle x \rangle$ broadcasts the message x to any recipient in transmission range.
- $\mathbf{r}(y)$ receives a message that has been broadcasted and binds it to y .

IV. METHODOLOGY

As previously discussed, autonomous vehicle platoons are a tantalizing goal due to the inherent benefits they bring in the form of increased road capacity and reduced energy usage [1], [2]. However, designing such a system is not a

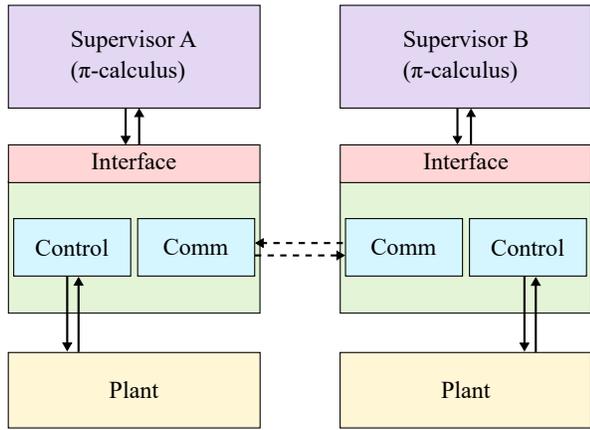


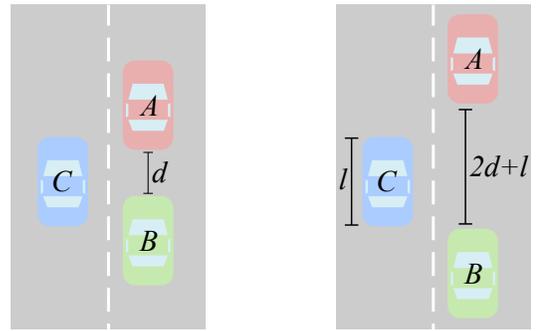
Fig. 1: Organization of two concurrent vehicle processes.

simple endeavor. In order to scale effectively over a large number of vehicles, this distributed platoon system should have decentralized control. Additionally, the benefits and stability of a platoon are increased if the vehicles remain in communication with each other [14]. As autonomous vehicles are safety-critical systems, the behavior that governs any such vehicle platoon must be verified to ensure no unforeseen dangerous scenarios can arise. Further complicating matters, autonomous vehicles are not homogeneous systems; each vehicle will have varying physical characteristics. Our work tackles these issues by dividing the problem into two non-overlapping parts: a high-level discrete-logic layer that describes the overall behavior of the autonomous vehicles, and a low-level physical layer that describes the closed-loop dynamics of the underlying system. Interaction occurs in the form of an interface provided by the low-level layer through which the high-level layer can send signals. This organization is shown in Fig. 1. The following sections describe these layers in more detail.

A. High-level Layer

The high-level layer must be capable of satisfying three conditions: it must be able to model the protocols of a distributed multi-vehicle system, it must be able to model communication between vehicles, and it must be amenable to verification processes for safety-critical systems. As it turns out, π -calculus meets these three conditions and is used to model the high-level behavior in this work. First, however, we must establish the expected platoon behavior.

Figure 2a depicts a simple platoon of two vehicles. The vehicle labeled A is the leader of this platoon and the vehicle labeled B is following A at a minimum safe distance d . A third vehicle, C , is outside of the platoon in a separate lane. Suppose C wishes to join the platoon and that the optimal place to do so is between A and B . Before C is permitted to join, the distance between A and B must be increased to $2d + l$, where l is the length of C as shown in Fig. 2b. This action ensures that the minimum safe distance is always respected. This behavior can be further simplified by letting C follow A at distance of d while B also follows C at a



(a) Initial A - B spacing (b) After space is created for C
Fig. 2: Expected spacing behavior of vehicle convoy.

distance of d , making explicit knowledge of l unnecessary.

In this work, we define separate behaviors for three different scenarios: a leader of a platoon (Leader), a follower in a platoon (Follower), and a vehicle joining a platoon (Joiner). These are the behaviors exhibited by vehicles A , B , and C , respectively, in Fig. 2. It is also possible for vehicles to transition between behaviors. For example, a vehicle joining a platoon will become a Follower once it has fully merged. At the highest level, the system modeled in Fig. 2 can be represented as a π -calculus expression. Let $A = Leader$, $B = Follower$, and $C = Joiner$. This system is then described by $A||B||C$. Each of these behaviors will now be described in terms of a π -calculus expression.

At the beginning of this section, it was stated that an interface layer is utilized by the behaviors to exert control over the underlying continuous system. The interface layer is defined as a concurrently executing π -calculus process. This process listens for messages over a predetermined set of communication channels, performs low-level processing, and in some cases replies with a message. Although this communication is considered atomic to π -calculus expressions, this will not be the case with respect to the low-level continuous layer. In particular, some actions take considerable time to execute and may wait until certain conditions unknown to the high-level layer are satisfied. The following communication channels are provided in the interface layer.

- *get_id* gets the identification of the vehicle. If B in Fig. 2a calls *get_id*, it will return B .
- *get_ldr* gets the leader of the vehicle. As an example, if B in Fig. 2a calls *get_ldr*, it will return A .
- *set_ldr* sets the leader of the vehicle to the vehicle with the given identification.
- *drive* tells the vehicle to drive forward with respect to the road geometry.
- *keep_dist* maintains a safe following distance d from the current leader.
- *check_join* checks whether the vehicle with the given identification occupies the position where this vehicle wants to join. For example, if C in Fig. 2a invokes *check_join* on B , it will return positive. In this case study, the positions at which vehicles join are predetermined for demonstration purposes. In future work, this

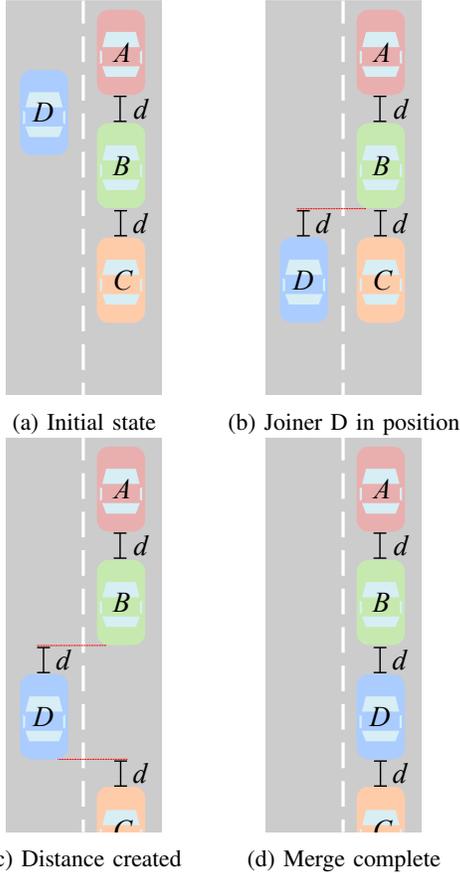


Fig. 3: Expected merge behavior of vehicle platoon.

will be replaced by an algorithm which decides where each vehicle should join.

- $align_start$ causes the interface to trigger the $align_done$ event action when the vehicle is d distance from the current leader.
- $merge_start$ causes the interface to trigger the $merge_done$ event action when the vehicle has finished merging into the current leader's lane.

Additionally, recursive expressions are only allowed to execute at most once per sampling period.

The behavior descriptions that follow refer to the system depicted in Fig. 3. A , B , and C are initially in a platoon of which A is the Leader and B and C are Followers. D is a Joiner external to the platoon.

1) *Leader (from [22])*: The behavior of a platoon Leader is straightforward in this scenario: it simply drives forward indefinitely. The associated process expression in Proc. 1 invokes the interface action $drive$ and repeats this action indefinitely. For the purposes of this paper, only straight roads are considered. However, this is easily extended to curved roads if lane-keeping is incorporated into the Leader.

Process 1 Leader

1: Leader = $\overline{drive}.Leader$

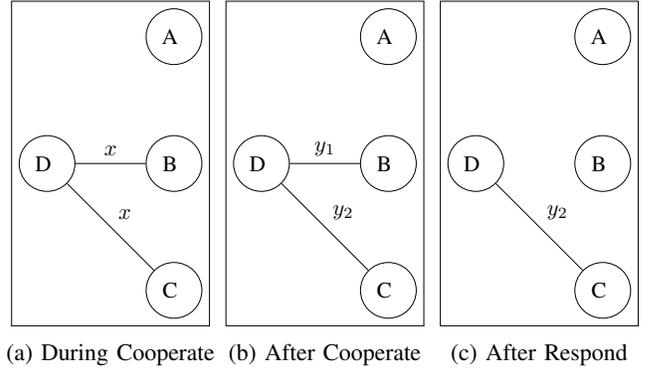


Fig. 4: Process graph depicting communication network during various points of execution. Cooperate and Respond refer to Lines 7 and 5 of Proc. 2 respectively. A , B , C , and D correspond to the vehicles in Fig. 3.

2) *Follower (from [22])*: The behavior of a Follower in a platoon is much more involved as it must simultaneously follow its leader while cooperating with vehicles who wish to join the platoon. This is modeled by the concurrent execution of Follow and Cooperate in the Follower state in Proc. 2, which corresponds with the initial state in Fig. 3a.

Process 2 Follower

- 1: Wait(y) = $\overline{y.merge_done}$
 - 2: Align(y) = $\overline{align_start}.align_done.\bar{y}.Wait(y)$
 - 3: Rcv_Ldr(y, ldr) = $y(nldr).set_ldr<nldr>.Align(y)$
 - 4: Send_Ldr(y) = $get_ldr(ldr).\bar{y}<ldr>.Rcv_Ldr(y, ldr)$
 - 5: Respond($y, flag$) = $flag : [True \Rightarrow Send_Ldr(y)]$
 - 6: Ident(y) = $get_id(id).\bar{y}<id>.y(flag).Respond(y, flag)$
 - 7: Cooperate = $\mathbf{!r}(x).\nu y(\bar{x}<y>).Ident(y)$
 - 8: Follow = $\overline{keep_dist}.Follow$
 - 9: Follower = Follow || Cooperate
-

The Cooperate expression demonstrates π -calculus's ability to dynamically reconfigure itself, which is an example of the mobility property. When a vehicle wishes to join a platoon, it creates a communication channel x and broadcasts it to every vehicle within range in Line 8 of Proc. 3. This establishes a common communication channel connecting the Joiner and every Follower. However, this is not a unique channel; any message sent on x can be received by every Follower. This is not ideal, so unique channels are created in Line 7 of Proc. 2, resulting in a reconfiguration of the network as shown in Fig. 4.

These communication channels are then used to synchronize the joining process. In Line 6 of Proc. 2, the Follower transmits its own identifier to the joining vehicle so that it can decide whether to join at the Follower's position. If this is the case, the Respond expression (Line 5) will continue to Send_Ldr and transmit the vehicle preceding the Follower to the Joiner. This enables the joining vehicle to position itself next to the Follower so that it is ready to merge over as shown in Fig. 3b. When the joining vehicle is in position,

the Follower sets the Joiner as its new preceding vehicle in Line 3 of Proc. 2. This causes the Follower to increase its distance to d from the joining vehicle due to the concurrently executing Follow expression. This is depicted in Fig. 3c. Lastly, the joining vehicle is signaled to merge, which results in Fig. 3d.

3) *Joiner*: The behavior of a Joiner is given in Proc. 3. Until a decision to join is made by the high-level layer, the joining vehicle drives forward with respect to the road geometry. Following this, the Joiner broadcasts its intention to join the platoon in Line 8 of Proc. 3. The interaction with Followers during the joining process has already been described; however, the decision of where to join in a platoon is handled by communication with the interface layer's *check_join* channel in the Check expression (Line 6). Once the Joiner has merged into the platoon, it transitions to a Follower as the last action in the Merge expression (Line 1).

Process 3 Joiner

```

1: Merge(y) =  $\overline{merge\_start.merge\_done}.\bar{y}.Follower$ 
2: Wait(y) =  $get\_id(id).\bar{y}<id>.y.Merge(y)$ 
3: Align(y) =  $\overline{align\_start.align\_done}.Wait(y)$ 
4: Rcv_Ldr(y) =  $y(ldr).set\_ldr < ldr > .Align(y)$ 
5: Ans(y, ok) =  $\bar{y}<ok>.ok : [True \Rightarrow Rcv\_Ldr(y)]$ 
6: Check(y, id) =  $(\nu z)(\overline{check\_join}<z, id>.z(ok)).Ans(y, ok)$ 
7: Listen(x) =  $x(y).y(id).Check(y, id)$ 
8: Joiner =  $(\nu x)(\bar{\mathbf{b}}<x>||!Listen(x))$ 

```

B. Low-level Layer

The low-level layer can be modeled as a series of hybrid automata [5] in order to describe the behavior of the controller and plant from Fig. 1. Each behavior – Leader, Follower, and Joiner – has its own hybrid automaton as shown in Fig. 5. Interaction with the high-level layer takes place in the form of input and output messages that directly correspond to communication channels provided by the interface. We now formally define these signals with respect to the hybrid automata. Let \mathbb{N} denote the set of natural numbers, then we define:

$$\begin{aligned}
S1 &: \{align_start, align_done\} \rightarrow \{absent, present\} \\
S2 &: \{merge_start, merge_done\} \rightarrow \{absent, present\} \\
S3 &: \{L, L'\} \rightarrow \mathbb{N}
\end{aligned}$$

Additionally, the following functions are defined. Let the x -axis lie parallel to the front of the vehicle and the y -axis lie perpendicular to that.

- $g_y(y)/g_x(x)$ computes \dot{y}/\dot{x} given the current y/x position while taking into account road geometry. In most cases, this is equivalent to driving forward at a constant velocity.
- $f_y(y, \alpha)/f_x(x, \alpha)$ computes \dot{y}/\dot{x} given the current y/x position and a distance α with the objective of reducing α to the minimum safe following distance d .

- $dist_x(t, L)$ returns the distance along the x -axis to vehicle L at time t .
- $dist_y(t, L)$ returns the distance along the y -axis to the center of the lane currently occupied by vehicle L at time t . Currently platoons are not allowed to switch lanes so that discontinuities in this distance are avoided.

With these signals and functions, we can define each of the vehicle behaviors.

1) *Leader*: The Leader automaton simply remains in a drive state. No transitions are supported and no input or output signals are accepted.

2) *Follower*: The Follower automaton requires the input signal *align_start* to create longitudinal space in front of the vehicle when a merge has been requested. This is accomplished by a change in leader L , which coincides with a *set_ldr* call in the high-level layer. However, only f_y respects the change in leader; f_x continues to be a function of the previous leader, L' . This is so the vehicle does not sway in its lane attempting to position itself laterally behind a vehicle that is changing lanes. When the *merge_done* signal is received, L is respected on both axes.

3) *Joiner*: Similarly, the Joiner automaton requires the input signal *align_start* to position itself behind the vehicle it will be following once the merge is completed. However, the x -axis does not respect its new leader, L , until the *merge_start* signal is received – indicating that the Follower vehicle is finished aligning itself to this vehicle.

In practice, these automata are implemented as digital controllers that receive input signals from the high-level layer via an interface as described in Section IV-A. The advantage of this setup is that the controllers and vehicle dynamics can be swapped out for any implementation as long as it fulfills the required interface. As an illustrative example, consider a scenario in which the high-level behavior described in the previous section has been adopted as an automobile industry standard. Each vehicle manufacturer can then supply a low-level layer for controlling its vehicle without having to re-verify the behavior protocol.

V. EXPERIMENTS

Simulation is an important step in the analysis and verification of autonomous vehicle platoons, as the closed-loop dynamics are a fundamental part of the system. Without simulation, the dynamics are not validated, and much of the problem is left unsolved. To that end, we have developed a simulation framework which ties together the high-level and low-level layers discussed in this paper and places them in control of automobiles in the Webots 8.3 robotics simulator [26]. A simulation environment, π -calculus behaviors, and vehicle control models are provided as inputs. The framework parses the behaviors and generates executable code which handles all necessary communication and interaction as defined by the π -calculus expressions. Two types of environments were tested: a homogeneous platoon consisting of vehicles with identical physical characteristics and controllers, and a heterogeneous platoon with varying physical

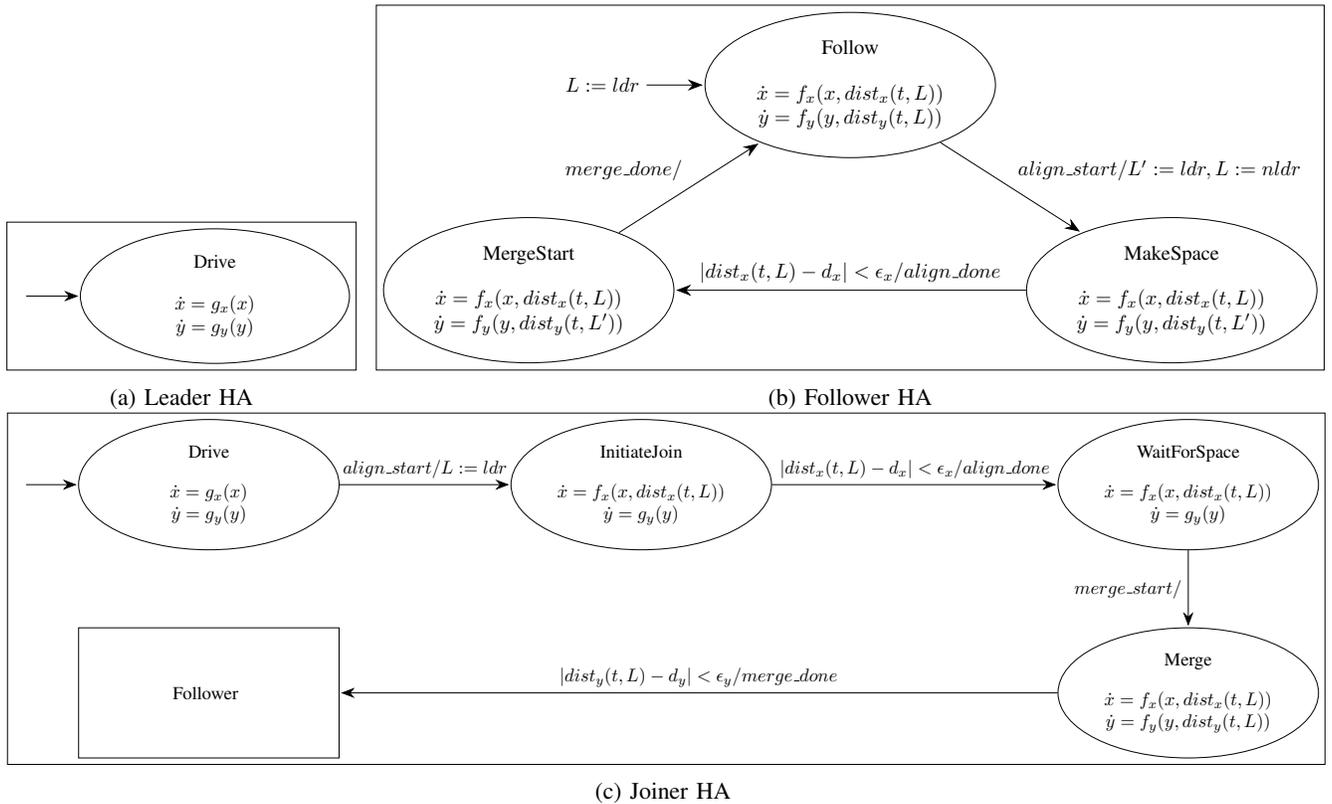


Fig. 5: Hybrid automata for the low-level layer. Each automaton receives input signals from the high-level layer and sends output signals. (a) No inputs/outputs. (b) Inputs: $align_start$, $merge_done$ Outputs: $align_done$ (c) Inputs: $align_start$, $merge_start$ Outputs: $align_done$, $merge_done$

characteristics and differing controllers. The behavior of vehicles in these simulations is shown in Fig. 6.

A. Homogeneous Platoon

In the first experimental simulation scenario, four independent vehicles attempted to join a five-vehicle platoon. All automobiles were modeled after the Toyota Prius (model included in Webots) with identical physical characteristics and identical manually tuned PID controllers. The controllers generate throttle and brake inputs for the plant (simulated vehicle) and track a target following distance from the vehicle which is being followed. The controllers are not meant to be robust nor optimal, rather they are a simple example of how this framework can be used to control continuous systems. Localization is achieved by means of each vehicle broadcasting out its current position at every simulation step. While not practical in the real world, this is adequate for this simulation as accurately identifying other vehicles and their positions is a challenging task. The distance functions, $dist_x$ and $dist_y$, make use of these broadcasted positions in addition to a forward-facing laser sensor for redundancy. All vehicles eventually successfully merge into the platoon in a distributed and decentralized manner.

B. Heterogeneous Platoon

The second experimental scenario again consisted of four independent vehicles attempting to join a five-vehicle pla-

toon; however, this time the vehicles were of varying models (included in Webots). Four automobiles were modeled after the BMW X5, three after the Citroen C-Zero, and two after the Toyota Prius. Additionally, three of the vehicles used the PID controller from the homogeneous platoon experiment while the remaining six used a model predictive controller designed for such scenarios with some minor modifications [27]. We vary the reference velocity based on the current tracking error, which is similar to previous work [28]. In addition we modify the cost parameters, Q and R of the controller, so that the weight given to the reference velocity and distance varies based on the error. Once again, all vehicles successfully merged into the platoon without incident.

However, similar to the PID controller in the first experimental simulation, this model predictive controller is not meant to be optimal nor even provide guarantees of stability. Rather its primary purpose is to show that our framework can not only effectively represent interactions between heterogeneous vehicles, but that those vehicles can use dramatically different control schemes.

VI. CONCLUSIONS

In this paper, we have introduced a new framework which utilizes π -calculus to model complex vehicular networks. This framework allows for the effective modeling of both

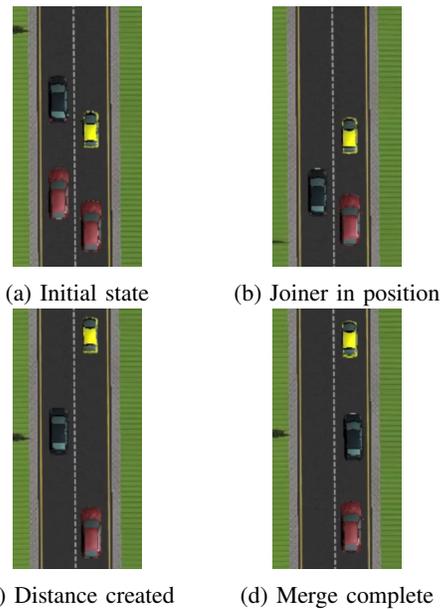


Fig. 6: Actual merge behavior of vehicle platoon.

high-level decision-making protocols and low-level vehicle dynamics, laying the groundwork for future verification of these systems. In order to demonstrate the utility of this framework, we have modeled a decentralized platooning protocol with heterogeneous vehicles and controllers. In future work, we will formally verify complex systems modeled with this framework.

REFERENCES

- [1] M. Zabat, N. Stabile, S. Farascarioli, and F. Browand, "The aerodynamic performance of platoons: A final report," California Partners for Advanced Transit and Highways (PATH), Tech. Rep. UCB-ITS-PRR-95-35, 1995.
- [2] B. Rao and P. Varaiya, "Flow benefits of autonomous intelligent cruise control in mixed manual and automated traffic," in *Intelligent Vehicle Highway Systems*. National Academy Press, 1993, no. 1408, pp. 36–43.
- [3] A. Colombo and D. Del Vecchio, "Efficient algorithms for collision avoidance at intersections," in *Proceedings of the 15th ACM International Conference on Hybrid Systems: Computation and Control*, 2012.
- [4] R. Milner, *Communicating and Mobile Systems: the π Calculus*. Cambridge University Press, 1999.
- [5] R. Alur, *Principles of Cyber-Physical Systems*. MIT Press, 2015.
- [6] S. Delaune, M. Ryan, and B. Smyth, "Automatic verification of privacy properties in the applied pi calculus," in *Trust Management II*, ser. IFIP, Y. Karabulut, J. Mitchell, P. Herrmann, and C. D. Jensen, Eds., vol. 263. Springer, 2008, pp. 263–278.
- [7] A. Singh, C. Ramakrishnan, and S. A. Smolka, "A process calculus for mobile ad hoc networks," *Science of Computer Programming*, vol. 75, no. 6, pp. 440–469, 2010.
- [8] G. Frehse, C. L. Guernic, A. Donz, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, and O. Maler, "Spaceex: Scalable verification of hybrid systems," in *Proceedings of the 23rd International Conference on Computer Aided Verification*, 2011.
- [9] Y. S. R. Annapureddy, C. Liu, G. E. Fainekos, and S. Sankaranarayanan, "S-TaLiRo: A tool for temporal logic falsification for hybrid systems," in *Tools and Algorithms for the Construction and Analysis of Systems*, ser. LNCS, vol. 6605. Springer, 2011, pp. 254–257.
- [10] J. Kim and I. Lee, "Modular code generation from hybrid automata based on data dependency," in *Real-Time and Embedded Technology and Applications Symposium, 2003. Proceedings. The 9th IEEE*. IEEE, 2003, pp. 160–168.
- [11] M. Althoff, D. Althoff, D. Wollherr, and M. Buss, "Safety verification of autonomous vehicles for coordinated evasive maneuvers," in *IEEE Intelligent Vehicles Symposium*, 2010.
- [12] W. Damm, A. Mikschl, J. Oehlerking, E.-R. Olderog, J. Pang, A. Platzer, M. Segelken, and B. Wirtz, "Automating verification of cooperation, control, and design in traffic applications," in *Formal Methods and Hybrid Real-Time Systems*, ser. LNCS, C. B. Jones, Z. Liu, and J. Woodcock, Eds. Springer, 2007, vol. 4700, pp. 115–169.
- [13] M. Asplund, A. Manzoor, M. Bourroche, S. Clarke, and V. Cahill, "A formal approach to autonomous vehicle coordination," in *FM 2012: Formal Methods*. Springer, 2012, pp. 52–67.
- [14] J. Lygeros, D. N. Godbole, and S. Sastry, "Verified hybrid controllers for automated vehicles," *IEEE Transactions on Automatic Control*, vol. 43, no. 4, pp. 522–539, 1998.
- [15] S. M. Loos, A. Platzer, and L. Nistor, "Adaptive cruise control: Hybrid, distributed, and now formally verified," in *FM 2011: Formal Methods*. Springer, 2011, pp. 42–56.
- [16] A. Platzer, "A complete axiomatization of quantified differential dynamic logic for distributed hybrid systems," *Logical Methods in Computer Science*, vol. 8, no. 4, 2012.
- [17] M. Fränzle, M. R. Hansen, and H. Ody, "No need knowing numerous neighbours," in *Correct System Design*. Springer, 2015, pp. 152–171.
- [18] M. Hilscher, S. Linker, E.-R. Olderog, and A. P. Ravn, "An abstract model for proving safety of multi-lane traffic manoeuvres," in *Formal Methods and Software Engineering*. Springer, 2011, pp. 404–419.
- [19] G. v. Bochmann, M. Hilscher, S. Linker, and E.-R. Olderog, "Synthesizing controllers for multi-lane traffic maneuvers," in *Dependable Software Engineering: Theories, Tools, and Applications*. Springer, 2015, pp. 71–86.
- [20] H. H. Bengtsson, L. Chen, A. Voronov, and C. Englund, "Interaction protocol for highway platoon merge," in *Proceedings of the 2015 IEEE 18th International Conference on Intelligent Transportation Systems*. IEEE, 2015, pp. 1971–1976.
- [21] M. Kamali, L. A. Dennis, O. McAree, M. Fisher, and S. M. Veres, "Formal verification of autonomous vehicle platooning," *arXiv preprint arXiv:1602.01718*, 2016.
- [22] J. Campbell, C. E. Tuncali, T. P. Pavlic, and G. Fainekos, "Toward modeling concurrency and reconfiguration in vehicular systems," in *9th Interaction and Concurrency Experience, Satellite Workshop of DisCoTec 2016*, 2016.
- [23] J. C. Baeten, "A brief history of process algebra," *Theoretical Computer Science*, vol. 335, no. 2, pp. 131–146, 2005.
- [24] R. Milner, J. Parrow, and D. Walker, "A calculus of mobile processes, I," *Information and Computation*, vol. 100, no. 1, pp. 1–40, 1992.
- [25] R. Milner, "The polyadic π -calculus: a tutorial," in *Logic and Algebra of Specification*, ser. NATO ASI Series, F. L. Bauer, W. Brauer, and H. Schwichtenberg, Eds. Springer, 1993, vol. 94, pp. 203–246.
- [26] Webots, "http://www.cyberbotics.com," commercial Mobile Robot Simulation Software. [Online]. Available: http://www.cyberbotics.com
- [27] P. Liu and U. Ozguner, "Predictive control of a vehicle convoy considering lane change behavior of the preceding vehicle," in *Proceedings of the 2015 American Control Conference*. IEEE, 2015, pp. 4374–4379.
- [28] D. N. Godbole and J. Lygeros, "Longitudinal control of the lead car of a platoon," *Vehicular Technology, IEEE Transactions on*, vol. 43, no. 4, pp. 1125–1135, 1994.