

Adaptive Cruise Control: Hybrid, Distributed, and Now Formally Verified*

Sarah M. Loos, André Platzer, and Ligia Nistor

Carnegie Mellon University, Computer Science Department, Pittsburgh, PA, USA
{sloos, aplatzer, lnistor}@cs.cmu.edu

Abstract. Car safety measures can be most effective when the cars on a street coordinate their control actions using distributed cooperative control. While each car optimizes its navigation planning locally to ensure the driver reaches his destination, all cars coordinate their actions in a distributed way in order to minimize the risk of safety hazards and collisions. These systems control the physical aspects of car movement using cyber technologies like local and remote sensor data and distributed V2V and V2I communication. They are thus cyber-physical systems. In this paper, we consider a distributed car control system that is inspired by the ambitions of the California PATH project, the CICAS system, SAFESPOT and PREVENT initiatives. We develop a formal model of a distributed car control system in which every car is controlled by adaptive cruise control. One of the major technical difficulties is that faithful models of distributed car control have both distributed systems and hybrid systems dynamics. They form distributed hybrid systems, which makes them very challenging for verification. In a formal proof system, we verify that the control model satisfies its main safety objective and guarantees collision freedom for arbitrarily many cars driving on a street, even if new cars enter the lane from on-ramps or multi-lane streets. The system we present is in many ways one of the most complicated cyber-physical systems that has ever been fully verified formally.

1 Introduction

Because of its societal relevance, many parts of car control have been studied before [1–18]. Major initiatives have been devoted to developing next generation individual ground transportation solutions, including the California PATH project, the SAFESPOT and PREVENT initiatives, and the CICAS-V system. Chang et al. [1], for instance, propose CICAS-V in response to a report that crashes at intersections in the US cost \$97 Billion in the year 2000. The promise is tempting. Current uncontrolled car traffic is inefficient and has too many safety risks, which are caused, e.g., by traffic jams behind curves, reduced vision at night, inappropriate reactions to difficult driving conditions, or sleepy drivers. Next generation car control aims to solve these problems by

* This material is based upon work supported by National Science Foundation under NSF CAREER Award CNS-1054246 and Grant Nos. CNS-0926181, CNS-0931985, CNS-1035800, CNS-1035813, and ONR N00014-10-1-0188. The first author was supported by an NSF Graduate Research Fellowship. For proofs and interactive car system simulations, see <http://www.ls.cs.cmu.edu/dccs/> online.

using advanced sensing, wireless V2V (vehicle to vehicle) and V2I (vehicle to roadside infrastructure) communication, and (semi)automatic driver assistance technology that prevents accidents and increases economical and ecological efficiency.

Yet, there are several challenges that still need to be solved to make next generation car control a reality. The most interesting challenge for us is that it only makes sense to introduce any of these systems after its correct functioning and reliability has been ensured. Otherwise, the system might do more harm than good. This is the formal verification problem for distributed car control, which we consider in this paper.

What makes this problem particularly exciting is its practical relevance. What makes it particularly challenging is its complicated dynamics. Distributed car control follows a hybrid dynamics, because cars move continuously along differential equations and their behavior is affected by discrete control decisions like when and how strongly to brake or to accelerate and to steer. It is in the very nature of distributed car control, however, to go beyond that with *distributed* traffic agents that interact by local sensing, broadcast communication, remote sensor data, or cooperative networked control decisions. This makes distributed car control systems prime examples of what are called *distributed hybrid systems*. In fact, because they form distributed cyber-physical multi-agent systems, the resulting systems are distributed hybrid systems regardless of whether they are built using explicitly distributed V2V and V2I network communication infrastructure or just rely on the distributed effects of sensor readings about objects traveling at remote locations (e.g., laser-range sensors measuring the distance to the car in front).

Cars reach maneuvering decisions locally in a distributed way. Is the global dynamics that emerges from the various local choices safe? What can a car assume about other cars in its maneuver planning? How do we ensure that multiple maneuvers that make sense locally do not cause conflicts or collisions globally? Formal verification of distributed hybrid systems had been an essentially unsolved challenge until recently [19].

Our main contribution is that we develop a distributed car control system and a formal proof that this system is collision-free for arbitrarily many cars, even when new cars enter or leave a multi-lane highway with arbitrarily many lanes. Another contribution is that we develop a proof structure that is strictly modular. We reduce the proof to modular stages that can be verified without the details in lower levels of abstraction. We believe the principles behind our modular structure and verification techniques are useful for other systems beyond the automotive domain. Further contributions are:

- This is the first case study in distributed hybrid systems to be verified with a generic and systematic verification approach that is not specific to the particular problem.
- We identify a simple invariant that all cars have to obey and show that it is sufficient for safety, even for emergent behavior of multiple distributed car maneuvers.
- We identify generic and static constraints on the input output parameters that any controller must obey to ensure that cars always stay safe.
- We demonstrate the feasibility of distributed hybrid systems verification.

2 Related Work

Car control is a deep area that has been studied by a number of different communities. The societal relevance of vehicle cooperation for CICAS intersection collision

avoidance [11] and for automated highway systems [5,8] has been emphasized. Horowitz et al. [10] proposed a lane change maneuver within platoons. Varaiya [13] outlines the key features of an IVHS (Intelligent Vehicle/Highway System). A significant amount of work has been done in the pioneering California PATH Project. Our work is strongly inspired by these systems, but it goes further and sets the groundwork for the modeling and formal verification of their reliability and safety even in distributed car control.

Dao et al. [3,4] developed an algorithm and model for lane assignment. Their simulations suggest [3] that traffic safety can be enhanced if vehicles are organized into platoons, as opposed to having random space between them. Our approach considers an even more general setting: we not only verify safety for platoon systems, but also when cars are driving on a lane without following platooning controllers. Hall et al. [6] also used simulations to find the best strategy of maximizing traffic throughput. Chee et al. [15] showed that lane change maneuvers can be achieved in automated highway systems using the signals available from on-board sensors. Jula et al. [9] used simulations to study the conditions under which accidents can be avoided during lane changes and merges. They have only tested safety partially. In contrast to [3,4,6,9,15], we do not use simulation but formal verification to validate our hypotheses.

Hsu et al. [7] propose a control system for IVHS that organizes traffic in platoons of closely spaced vehicles. They specify this system by interacting finite state machines. Those cannot represent the actual continuous movement of the cars. We use differential equations to model the continuous dynamics of the vehicles and thus consider more realistic models of the interactions between vehicles, their control, and their movement.

Stursberg et al. [12] applied counterexample-guided verification to a cruise control system with two cars on one lane. Their technique can not scale to an arbitrary number of cars. Althoff et al. [17] use reachability analysis to prove the safety of evasive maneuvers with constant velocity. They verify a very specific situation: a wrong way driver threatens two autonomously driving vehicles on a road with three lanes.

Wongpiromsarn et al. [14] verify safety of the planner-controller subsystem of a single autonomous ground vehicle. Their verification techniques restrict acceleration changes to fixed and perfect polling frequency, while our model of an arbitrary number of cars allows changes in acceleration at any point in time, with irregular sensor updates.

Damm et al. [2] give a verification rule that is specialized to collision freedom of traffic agents. To show that two cars do not collide, they need to manually prove eighteen verification conditions. Lygeros and Lynch [20] prove safety only for one deceleration strategy for a string of vehicles: the leading vehicle applies maximum deceleration until it stops, while at the same time, the cars following it in the string decelerate to a stop. The instantaneous, globally synchronized reaction of the cars is an unrealistic assumption that we do not make in our case study. Dolginova and Lynch [21] verify that no collisions with big relative velocity can occur when two adjacent platoons do a merge maneuver. This does not prove the absence of small relative velocity collisions, nor the behavior of 3 platoons or when not merging. In contrast to the manual semantic reasoning of [2,20,21], our techniques follow a formal proof calculus [19], which can be mechanized. In the case studies analyzed by [20,21] safety is proved only for a particular scenario, while our modular formal proofs deal with the general case. In our case study, the cars have more flexibility and an arbitrary number of control choices.

Unlike [2,12,14,17], we prove safety for an arbitrary number of cars. Our techniques and results are more general than the case-specific approaches [2,12,14,17,20,21], as we prove collision-freedom for any number of cars driving on any finite number of lanes. None of the previously cited papers have proved safety for distributed car control in which cars can dynamically enter the highway system, change lanes, and exit.

3 Preliminaries: Quantified Differential Dynamic Logic

Distributed car control systems are distributed hybrid systems, which we model by *quantified hybrid programs* (QHPs) [19]. QHPs are defined by the grammar (α, β are QHPs, θ a term, i a variable, f a function symbol, and H a formula of first-order logic):

$$\alpha, \beta ::= \forall i : C f(i) := \theta \mid \forall i : C f(i)' = \theta \& H \mid f(i) := * \mid ?H \mid \alpha \cup \beta \mid \alpha ; \beta \mid \alpha^*$$

The effect of *quantified assignment* $\forall i : C f(i) := \theta$ is an instantaneous discrete jump assigning θ to $f(i)$ simultaneously for all objects i of type C . Usually i occurs in θ . The effect of *quantified differential equation* $\forall i : C f(i)' = \theta \& H$ is a continuous evolution where, for all objects i of type C , all differential equations $f(i)' = \theta$ hold *and* (written $\&$ for clarity) formula H holds throughout the evolution (the state remains in the region described by H). Usually, i occurs in θ . Here $f(i)'$ is intended to denote the derivative of the interpretation of the term $f(i)$ over time during continuous evolution, not the derivative of $f(i)$ by its argument i . For $f(i)'$ to be defined, we assume f is an \mathbb{R} -valued function symbol. The effect of the random assignment $f(i) := *$ is to non-deterministically pick an arbitrary number or object (of type the type of $f(i)$) as the value of $f(i)$.

The effect of *test* $?H$ is a *skip* (i.e., no change) if formula H is true in the current state and *abort* (blocking the system run by a failed assertion), otherwise. *Non-deterministic choice* $\alpha \cup \beta$ is for alternatives in the behavior of the distributed hybrid system. In the *sequential composition* $\alpha ; \beta$, QHP β starts after α finishes (β never starts if α continues indefinitely). *Non-deterministic repetition* α^* repeats α an arbitrary number of times ≥ 0 .

For stating and proving properties of QHPs, we use *quantified differential dynamic logic* QdL [19] with the grammar:

$$\phi, \psi ::= \theta_1 = \theta_2 \mid \theta_1 \geq \theta_2 \mid \neg\phi \mid \phi \wedge \psi \mid \forall i : C \phi \mid \exists i : C \phi \mid [\alpha]\phi \mid \langle \alpha \rangle \phi$$

In addition to all formulas of first-order real arithmetic, QdL allows formulas of the form $[\alpha]\phi$ with a QHP α and a formula ϕ . Formula $[\alpha]\phi$ is true in a state ν iff ϕ is true in all states that are reachable from ν by following the transitions of α ; see [19] for details.

4 The Distributed Car Control Problem

Our approach to proving safety of a distributed car control system is to break the verification into modular pieces. In this way, we simplify what would otherwise be a very large and complex proof. The ultimate result of this paper is a formally verified model of any straight stretch of highway on which each car is following adaptive cruise control. On any highway, there will be an arbitrary number of lanes and an arbitrary number of cars, and the system will change while it runs when cars enter and leave the highway.

This would be an incredibly complex system to verify if we were to tackle it at this level. Each lane has a group of cars driving on it. This group is constantly changing as cars weave in and out of surrounding traffic. Each car has a position, velocity, and acceleration, and must obey the laws of physics. On top of that, in order to ensure complete safety of the system, every car must be certain at all times that its control choices will not cause a collision anywhere else in the system at any time in the future.

These issues are compounded by the limits of the sensory and communications networks. On a highway that stretches hundreds of miles, we could not hope for any car to collect and analyze real-time data from every other car on the interstate. Instead, we must assume each car is making decisions based on its local environment, e.g., within the limitations of sensors, V2V and V2I communication, and real-time computation.




Fig. 1. Emergent highway collision risk

Additionally, once you split your system into reasonably local models, it is still difficult to reason about how these local groups of cars interact. For example, consider a local group of three cars for a lane change maneuver: the car changing lanes, and the two cars that will be ahead and behind it. It is tempting to signal the car ahead to speed up and the car behind to slow down in

order to make space for the car changing lanes. This is perfectly reasonable on the local level; however, Fig. 1 demonstrates a problem that appears when we attempt to compose these seemingly safe local cases into a global system. Two cars are attempting safe and legal lane changes simultaneously, but the car which separates the merging cars is at risk. The car in the middle simultaneously receives requests to slow down and speed up. It cannot comply, which could jeopardize the safety of the entire system.

To avoid complex rippling cases that could result in a situation similar to the one in Fig. 1, we organize our system model as a collection of hierarchical modular pieces. The smallest piece consists of only two cars on a single lane. We present a verification of this model in Sect. 5 and build more complex proofs upon it throughout the paper.

In Sect. 6, we prove that a lane with an arbitrary number of cars driven by any distributed homogeneous adaptive cruise control system is safe, assuming the system has been proved safe for two cars. We generate our own verified adaptive cruise control model for this system, but, due to the modular proof structure, it can be substituted with *any* implementation-specific control system which has been proved safe for two cars.

The verification of this one lane system, as well as the verification we present in Sect. 8 for a highway with multiple lanes, will hold independently with respect to the adaptive cruise control specifications. In Sect. 7, we look at the local level of a multi-lane highway system. We verify the adaptive cruise control for a single lane, where cars are allowed to merge in and out of the lane. Finally in Sect. 8, we compose the lane systems verified in Sect. 7 to provide a full verification of the highway system.

5 Local Lane Control

The local car dynamics problem that we are solving is: we have two cars on a straight lane that can accelerate, coast or brake and we want to prove that they will not collide.

This system contains complex physical controls as well as discrete and continuous dynamics, thus, is a hybrid system. Once the model for the local problem is verified, we will use it in a compositional fashion to prove safety for more complicated scenarios, such as multiple cars driving on a lane or on parallel lanes. We can apply modular composition because we have structured the models in a hierarchical order, we have found the right decomposition of the sub-problems and we have identified the right invariants.

Modeling. We develop a formal model of the local car dynamics as a QHP. Each car has state variables that determine how it operates: position, velocity, and acceleration. For follower car f , x_f represents its position, v_f its velocity, and a_f its acceleration (similarly for leader car ℓ).

The continuous dynamics for f are described by the following differential equation system: $x'_f = v_f$, $v'_f = a_f$. This is the ideal-world dynamics that is adequate for a kinematic model of longitudinal lane maneuvers. The rate with which the position of the car changes is given by x'_f , i.e., the velocity. The velocity itself changes continuously according to the current acceleration a_f . We do not assume permanent control over the acceleration, but tolerate delays since sensor readings are not available continuously, control decisions may need time, and actuators may take time to react. For simplicity, though, we still assume that, once set, the acceleration a_f takes instant effect. We assume a global limit for the maximum acceleration and we denote it by $A \geq 0$. We assume that all cars have an emergency brake with a braking power between a maximum value B and a minimum value b , where $B \geq b > 0$. The two values have to be positive, otherwise the cars cannot brake. They may be different, however, because we cannot expect all cars to realize exactly the same emergency braking power and it would be unrealistic to build a system based on the assumption that all reactions are equal.

In Fig. 2, we see that leader ℓ brakes unexpectedly at time t_1 with its maximum braking power, $-B$. Unfortunately, f did not follow ℓ at a safe distance, and so when sensor and network data finally inform f at time t_2 that ℓ is braking, it is already too late for f to prevent a collision. Although f applies its full braking power, $-b$, at time t_2 , the cars will inevitably crash at time t_3 . The same problem can happen if ℓ brakes with $-b$ and f brakes with $-B$. This example shows that control choices which look good early on can cause problems later. Adding cars to the system amplifies these errors.

We present the entire specification of the local lane control (11c), consisting of the discrete control and the continuous dynamics, in Model 1. This system evolves over time, which is measured by a clock, i.e., variable t changing with slope $t' = 1$ as in (8). The differential equation system (8) formalizes the physical laws for movement, which are restricted to the evolution domain (9). Neither human drivers nor driver assistance




Fig. 2. Local car crash

Model 1. Local lane control (11c)

$$11c \equiv (ctrl; dyn)^* \quad (1)$$

$$ctrl \equiv \ell_{ctrl} \parallel f_{ctrl}; \quad (2)$$

$$\ell_{ctrl} \equiv (a_\ell := *; ?(-B \leq a_\ell \leq A)) \quad (3)$$

$$f_{ctrl} \equiv (a_f := *; ?(-B \leq a_f \leq -b)) \quad (4)$$

$$\cup (?Safe_\varepsilon; a_f := *; ?(-B \leq a_f \leq A)) \quad (5)$$

$$\cup (?(v_f = 0); a_f := 0) \quad (6)$$

$$Safe_\varepsilon \equiv x_f + \frac{v_f^2}{2b} + \left(\frac{A}{b} + 1\right) \left(\frac{A}{2}\varepsilon^2 + \varepsilon v_f\right) < x_\ell + \frac{v_\ell^2}{2B} \quad (7)$$

$$dyn \equiv (t := 0; x'_f = v_f, v'_f = a_f, x'_\ell = v_\ell, v'_\ell = a_\ell, t' = 1) \quad (8)$$

$$v_f \geq 0 \wedge v_\ell \geq 0 \wedge t \leq \varepsilon \quad (9)$$

technology are able to react immediately and each vehicle or driver will have a specific reaction time. Therefore we have a constant parameter, ε , which serves as an upper bound on the reaction time for all vehicles. We verify car control for arbitrary values of ε . Cars can react as quickly as they want, but they can take no longer than ε .

The leading car is not restricted by the car behind, so it may accelerate, coast, or brake at will. In Model 1, a_ℓ is first randomly assigned a real value, non-deterministically through (3). The model continues if a_ℓ is within the physical limits of the car's brakes and engine, i.e. between $-B$ and A . On the other hand, f depends on the distance to ℓ and has a more restrictive set of possible moves. Car f can take some choices only if certain safety constraints about the distance and velocities are met.

Braking is allowed at all times, so a human driver may always override the automated control to brake in an emergency. In fact, braking is the only option if there is not enough distance between the cars for f to maintain its speed or accelerate. This is represented in (4), where there is no precondition for any force between $-B$ and $-b$.

The second possibility, (5), is that there is enough distance between the two cars for f to take any choice. This freedom is only given when (7) is satisfied. If (7) holds, then ℓ will still be safely in front of f until the controllers can react again (i.e., after they drive for up to ε time units), no matter how ℓ accelerates or brakes. This distance is greater than the minimum distance required for safety if they both brake simultaneously. The ε terms in (7) add this extra distance to account for the possibility that f accelerates for time ε even when ℓ decides to brake, which f may not notice until the next sensor update. These terms represent the distance traveled during one maximum reaction cycle of ε time units with worst-case acceleration A , including the additional distance needed to reduce the speed down to v_f again after accelerating with A for ε time units.

Now the third possibility. If f had previously chosen to brake by $a_f = -b$ then the continuous evolution dyn cannot continue with the current acceleration choices below velocity $v_f = 0$ due to constraint (9). Thus, we add the choice (6) saying that the car may always choose to stand still at its position if its velocity is 0 already.

The two cars can repeatedly choose from the range of legal accelerations. This non-deterministic repetition is represented by operator $*$ in (1). The controllers of the two cars operate in parallel as seen in (2). Notice that the controllers are independent with

respect to read and write variables (which also makes sense for implementation purposes), so in this case, parallel (\parallel) is equivalent to sequential composition ($;$).

Verification. To verify the local lane control problem modeled in Sect. 5, we use a formal proof calculus for QdL [19]. In the local lane control problem, we want f to be safely behind ℓ at all times. To verify that a collision is not possible, we show that there is always a reasonable distance between ℓ and f ; enough distance that if both cars brake instantly, the cars would not collide. We verify this property for all times and under any condition which the system can run, so if a car can come so close to another car that even instant braking would not prevent a crash, the system is already unsafe.

For two cars f and ℓ , we have identified the following crucial relation ($f \ll \ell$), i.e., follower f is *safely behind* leader ℓ :

$$(f \ll \ell) \equiv (x_f \leq x_\ell) \wedge (f \neq \ell) \rightarrow \left(x_f < x_\ell \wedge x_f + \frac{v_f^2}{2b} < x_\ell + \frac{v_\ell^2}{2B} \wedge v_f \geq 0 \wedge v_\ell \geq 0 \right)$$

If ($f \ll \ell$) is satisfied, then f has a safe distance from ℓ . The formula states that, if ℓ is the leading car (i.e., $x_f \leq x_\ell$ for different cars $f \neq \ell$), then the leader must be strictly ahead of the follower, and there must be enough distance between them such that the follower can stop when the leader is braking. Also both cars must be driving forward.

The safe distance formula ($f \ll \ell$) is the most important invariant. The system must satisfy it at all times to be verified. This is not to be confused with the definition of Safe_ε in the control, which must foresee the impact of control decisions for the future of ε time. For simplicity, these formulas do not allow cars to have non-zero length; however, adding the car length to x_f would eliminate this requirement.

Proposition 1 (Safety of local lane control 11c). *If car f is safely behind car ℓ initially, then the cars will never collide while they follow the 11c control model; therefore, safety of 11c is expressed by the provable formula: $(f \ll \ell) \rightarrow [11c](f \ll \ell)$*

We proved Proposition 1 using KeYmaera, a theorem prover for hybrid systems (proof files available online [22]). A proof sketch is presented in [23, Appendix A.1].

6 Global Lane Control



Fig. 3. Lane risk

In Sect. 5 we show that a system of two cars is safe, which gives a local version of the problem to build upon. However, our goal is to prove safety for a whole highway of high-speed vehicles. The next step toward this goal is to verify safety for a single lane of n cars, where n is arbitrary and finite, and the ordering of the cars is fixed (i.e., no car can pass another). Each car follows the same control we proved safe for two cars in Sect. 5, but adding cars to the system and making it distributed has introduced new risks. It is now necessary to show, for example, if you are driving along and the car in front of you slows while the car behind simultaneously accelerates, you won't be left sandwiched between with no way to avoid a collision (as in Fig. 3).

Model 2. Global lane control (g1c)

$$\text{g1c} \equiv (\text{ctrl}^n; \text{dyn}^n)^* \quad (10)$$

$$\text{ctrl}^n \equiv \forall i : C (\text{ctrl}(i)) \quad (11)$$

$$\text{ctrl}(i) \equiv (a(i) := *; ?(-B \leq a(i) \leq -b)) \quad (12)$$

$$\cup (? \text{Safe}_\varepsilon(i); a(i) := *; ?(-B \leq a(i) \leq A)) \quad (13)$$

$$\cup (?(v(i) = 0); a(i) := 0) \quad (14)$$

$$\text{Safe}_\varepsilon(i) \equiv x(i) + \frac{v(i)^2}{2b} + \left(\frac{A}{b} + 1\right) \left(\frac{A}{2} \varepsilon^2 + \varepsilon v(i)\right) < x(L(i)) + \frac{v(L(i))^2}{2B} \quad (15)$$

$$\text{dyn}^n \equiv (t := 0; \forall i : C (\text{dyn}(i)), t' = 1, t \leq \varepsilon) \quad (16)$$

$$\text{dyn}(i) \equiv x'(i) = v(i), v'(i) = a(i), v(i) \geq 0 \quad (17)$$

Modeling. Because we are now looking at a lane of cars, our model will require additional features. First, we must represent the position, velocity, and acceleration of each car. If these variables were represented as primitives, the number of variables would be large and difficult to handle. Using only primitive variables, we cannot verify a system for any arbitrary number of cars, i.e., we could verify for, say, 5 cars, but not for any n cars. Therefore, we give each car an index, i , and use first-order variables $x(i)$, $v(i)$, and $a(i)$ to refer to the position, velocity and acceleration of car i . With these first-order variables, our verification applies to a lane of any number of cars.

Of course, the cars are all driving along the road at the same time, so we evolve the positions of the cars simultaneously along their differential equations. The acceleration, $a(i)$, of all cars is also set simultaneously in the control. We need notation for this parallel execution, so we use the universal quantifier (\forall) in the definition of the overall control and continuous dynamics (see (11) and (16) in Model 2). The control of all cars in the system is defined by ctrl^n (11). This says that for each car i , we execute $\text{ctrl}(i)$. This control is exactly the control defined in Sect. 5 - under *any* conditions the car may brake (12); if the car is safely following its leader, it may choose any valid acceleration between $-b$ and A (13); and if the car is stopped, it may remain stopped (14). There are only two distinctions between the control introduced in g1c and the control used in 11c described in Sect. 5. First, we change primitive variables to first-order variables. Second, with so many cars in the system, we have to determine which car is our leader.

It is vital that every car be able to identify, through local sensors or V2V/V2I communication networks, which car is directly in front of it. It is already assumed that the sensor and communication network is guaranteed to give accurate updates to every car within time ε . We now also make the reasonable assumption that with each update, every car is able to identify which car is directly ahead of it in its lane. This may be a bit tricky if the car only has sensor readings to guide it, but this assumption is reasonable if all cars are broadcasting their positions (and which lane they occupy in the case of multiple lanes). For some car i , we call the car directly ahead of it $L(i)$, or the *leader of car i* . More formally, we assume the following properties about $L(i)$:

$$L(i) = j \equiv x(i) < x(j) \wedge \forall k : C \setminus \{i, j\} (x(k) < x(i) \vee x(j) < x(k))$$

$$(i \ll L(i)) \equiv \forall j : C ((L(i) = j) \rightarrow (i \ll j))$$

The equation $L(i) = j$ is expanded to mean that the position of j must be ahead of the position of i , and there can be no cars between. The second formula states that for a car, i , to be safely behind its leader, denoted ($i \ll L(i)$), we require that i should be safely behind any car which fulfills the requirements of the first equation. Each car will have at most one leader at any given time. At the end of the finite length lane, we position a stationary car. This car has no leader and therefore will never move.

The constraint **Safe _{ε}** from Sect. 5 has been updated to a first-order variable as well (15). It now uses $L(i)$ to identify which car is directly ahead of car i , and then determines if i is following safely enough to accelerate for ε time. This constraint is applied to all cars in the system when the individual controls set acceleration.

The continuous dynamics are the same as those described in Sect. 5, but with the added dynamics of the other cars in the system (16). Once $a(i)$ has been set for all cars by ctr^m (11), each car evolves along the dynamics of the system for no more than ε time (maximum reaction time). The position of each car evolves as the second derivative of the acceleration set by the control (17). The model requires that the cars never move backward by adding the constraint $v(i) \geq 0$. We still have a global time variable, t , that is introduced in the definition of dyn^n (16). Since $t' = 1$, all cars evolve along their respective differential equations in an absolute timeframe. Note that t is never read by the controller, thus, g1c has no issues with local clock drift.

We model all cars in the system as repeatedly setting their accelerations as they synchronously receive sensor updates (11) and following the continuous dynamics (16). When put together and repeated non-deterministically with the $*$ operator, these QHPs form the g1c model (10) for global lane control. The g1c model is easy to implement since each car relies on local information about the car directly ahead. Our online supplementary material shows a demo of an implementation of this model [22].

Verification. Now that we have a suitable model for a system of n cars in a single lane, we identify a suitable set of requirements and prove that our model never violates them. In Sect. 5, since there were only two cars on the road, it was sufficient to show that the follower car was safely behind its leader at all times. However, in this model it is not enough to only ensure safety for each car and its direct leader. We must also verify that a car is safely following all cars ahead – each car has to be safely behind its leader, and the leader of its leader, and the car in front of that car, and so on.

For example, suppose there were a long line of cars following each other very closely (they could, for instance, be in a platoon). If the first car brakes, then one-by-one the cars behind each react to the car directly in front of them and apply their brakes. In some models, it would be possible for these reaction delays to add up and eventually result in a crash [24]. Our model is not prone to this fatal error, because our controllers are explicitly designed to tolerate reaction delays. Each car is able to come to a full stop no matter what the behavior of the cars in front of it (so long as all cars behave within the physical limits of their engines and brakes). To show this, we must verify that under the system controls every car is always safely behind all cars ahead until the lane ends. We do this by first defining *transitive leaders*, $L^*(i)$ as follows:

$$(i \ll L^*(i)) \equiv [k := i; (k := L(k))^*](i \ll k)$$

The QHP, $k := i; (k := L(k))^*$, continually redefines k to be the next car in the lane (until the lane ends). Because this QHP is encapsulated in [], all states that are

reachable in the program must satisfy the formula ($i \ll k$). In other words, starting with ($k := i$), we check that i is safely behind k , or ($i \ll i$). Next, $k := L(k)$, so $k := L(i)$, and we prove that i is safely behind k : ($i \ll L(i)$). Then we redefine k to be its leader again ($k := L(k)$), and we check that i is safely behind k : ($i \ll L(L(i))$). This check is continued indefinitely: ($i \ll L(L(\dots L(i)))$). Hence the notation, ($i \ll L^*(i)$).

Proposition 2 (Safety of global lane control g1c). *For every configuration of cars in which each car is safely following the car directly in front of it, all cars will remain in a safe configuration (i.e., no car will ever collide with another car) while they follow the distributed control. This is expressed by the following provable formula:*

$$\forall i : C(i \ll L(i)) \rightarrow [\text{g1c}](\forall i : C(i \ll L^*(i)))$$

This means that as the cars move along the lane, every car in the system is safely following all of its transitive leaders.

Using Gödel’s generalization rule, our proof for a lane of cars splits immediately into two branches: one which relies on the verification of the control and dynamics in the local, two car case, and one which verifies the rest of the system. These two branches are independent, and furthermore, the control and dynamics of the cars are only expanded in the verification of the local model. This is good news for two reasons. First, it keeps the resulting proof modular, which makes it possible to verify larger and more complex systems. Second, if the control or dynamics of the model are modified, only an updated verification of safety for two cars will be needed to verify the new model for the whole system. Proof details are available in [23, Appendix A.2].

7 Local Highway Control

In Sect. 6, we verified an automated control system for an arbitrary, but constant, number of cars on a lane. Later, we will put lots of these lanes together to model highway traffic. In our full highway model, cars will be able to pass each other, change lanes, and enter or leave the highway. We first study how this full system behaves from the perspective of a single lane. When a car changes into or out of that lane, it will look like a car is appearing or disappearing in the middle of the lane: in front of and behind existing cars. It is crucial to show that these appearances and disappearances are safe.

If a new car cuts into the lane without leaving enough space for the car behind it, it could cause an accident. Furthermore, when two cars enter the lane simultaneously, if there are several cars between them, we must prove that there will not be a ripple effect which causes those cars between to crash (also see Fig. 1). Faithful verification must apply to all kinds of complex maneuvers and show safety for all cars in the system, not just those involved locally in one maneuver.

Our verification approach proves separate, modular properties. This allows us to compose these modular proofs and verify collision freedom for the entire system for any valid maneuver, no matter how complex, even multiple maneuvers at different places.

Modeling. We have additional challenges in modeling this new system where cars can appear and disappear dynamically. First of all, in previous sections we have used $\forall i : C$ to mean “for all cars in the system.” We will now abuse this notation and take it to mean

Model 3. Local highway control (lhc)

$$\text{lhc} \equiv (\text{delete}^*, \text{create}^*, \text{ctrl}^n; \text{dyn}^n)^* \quad (18)$$

$$\text{create} \equiv n := \text{new}; ?((F(n) \ll n) \wedge (n \ll L(n))) \quad (19)$$

$$(n := \text{new}) \equiv n := *; ?(E(n) = 0); E(n) := 1 \quad (20)$$

$$(F(n) \ll n) \equiv \forall j : C (L(j) = n \rightarrow (j \ll n)) \quad (21)$$

$$\text{delete} \equiv n := *; ?(E(n) = 1); E(n) := 0 \quad (22)$$

“for all cars which currently exist on this lane.” (In our formal proof we use an actualist quantifier to distinguish between these situations. This technique is described in detail in another paper [19].) Secondly, our model must represent what physical conditions in the lane must be met before a car may disappear or appear safely. And finally, the model must be robust enough to allow disappearances and appearances to happen throughout the evolution of the system (i.e., a car may enter or leave the lane at any time).

Recall that a car, n , has three real values: position, velocity and acceleration. Now that cars can appear and disappear, we add a fourth element: existence. The existence field is just a bit that we flip on ($E(n) := 1$) when the car appears and flip off ($E(n) := 0$) when the car disappears.

When we create a new car, n , we start by allowing the car to be anything. This can be written in dynamic logic as a random assignment $n := *$. Of course, when we look at the highway system as a whole, we won't allow cars to pop out of thin air onto the lane. This definition can be restricted to cars which already exist on an adjacent lane. However, since the choice of $*$ is non-deterministic, we are verifying our model for all possible values of n . This means that the verification required for an entire highway system will be a subset of the cases covered by this model of a single lane. Because $n := *$ allows n to be any car, one that exists on the lane or one that doesn't, we first must check that this “new” car isn't already on the lane. If it doesn't exist, i.e. $?(E(n) = 0)$, then we can flip our existence bit to on and it will join the existing cars on this lane (20).

Now that we have defined appearance, we can define its dual: disappearance. We delete cars by choosing a car, n , non-deterministically, checking that it exists, and then flipping that bit so that it no longer exists on this lane (22). After a delete, notice that while the car ceases to exist physically on our lane, we are still able to refer to it in our model and verification as car n – a car that used to be in the lane.

A car may leave the lane at any time (assuming there is an adjacent lane which it can move into safely), but it should only be allowed to enter the lane if it is safely between the car that will be in front of it and the car that will be behind it. Because of this, when creating a car in the lane, our model will check that the car is safely between the car in front and behind. If we have a test which follows a creation of a new car, as in our definition of *create* in (19), a new car will only appear if the test succeeds. The formula $(F(i) \ll i)$ evaluates to true if car i is safely ahead of the car behind it. This is the dual of $(i \ll L(i))$. We define this formally in terms of $(i \ll L(i))$ as shown in (21).

The lhc model is identical to the glc model in Sect. 6, but at the beginning of each control cycle it includes zero or more car *deletes* or *creates* as shown by *delete*^{*} and *create*^{*} in (18). It is important to note that the verification will include interleaving and

simultaneous *creates* and *deletes* since the continuous dynamics (dyn^n) are allowed to evolve for zero time and start over immediately with another *delete* and *create* cycle.

Verification. Now that we have a model for local highway control, we have to describe a set of requirements that we want the model to satisfy in order to ensure safety. These requirements will be identical to the requirements necessary in the global lane control. We want to show that every car is a safe distance from its transitive leaders: $\forall i : C(i \ll L^*(i))$. Because these requirements are identical to those presented in Proposition 2, the statement of Proposition 3 is identical except for the updated model.

Proposition 3 (Safety of local highway control lhc). *Assuming the cars start in a controllable state (i.e. each car is a safe distance from the car in front of it), the cars may move, appear, and disappear as described in the (lhc) model, then no cars will ever collide. This is expressed by the following provable formula:*

$$\forall i : C(i \ll L(i)) \rightarrow [\text{lhc}]\forall i : C(i \ll L^*(i))$$

We keep the proof of Proposition 3 entirely modular just as we did in the previous section for Proposition 2. The proof is presented in [23, Appendix A.3].

8 Global Highway Control

So far, we have verified an automated car control system for cars driving on one lane. A highway consists of multiple lanes, and cars may change from one lane to the other. Just because a system is safe on one lane does not mean that it would operate safely on multiple lanes. When a car changes lanes, it might change from a position that used to be safe for its previous lane over to another lane where that position becomes unsafe. Lane change needs to be coordinated and not chaotic. We have to ensure that multiple local maneuvers cannot cause global inconsistencies and follow-up crashes; see Fig. 1.

Modeling. The first aspect we need to model is which lane is concerned. The quantifier $\forall i : C$, which in Sect. 7 quantified over “all cars which exist on the lane”, now needs to be parametrized by the lane that it is referring to. We use the notation $\forall i : C_l$ to quantify over all cars on lane l . Likewise, instead of the existence function $E(i)$, we now use $E(i, l)$ to say whether car i exists on lane l . A car could exist on some l but not on others. A car can exist on multiple lanes at once if its wheels are on different lanes (e.g., when crossing dashed lines). We use subscripted $ctr_l^m, dyn_l^n, L_l(i), L_l^*(i)$ etc. to denote variants of $ctr^m, dyn^n, L(i), L^*(i)$ in which all quantifiers refer to lane l . Similarly, we write $\forall l : L ctr_l^m$ for the QHP running the controllers of all cars on all lanes at once.

In addition to whatever a car may do in terms of speeding up or slowing down, lane change corresponds to a sequence of changes in existence function $E(i, l)$. A model for an instant switch of car i from lane l to lane l' would correspond to $E(i, l) := 0$; $E(i, l') := 1$, i.e., disappearance from l and subsequent appearance on l' . This is mostly for adjacent lanes $l' = l \pm 1$, but we allow arbitrary lanes l, l' to capture highways with complex topology. Real cars do not change lanes instantly, of course. They gradually move from one lane over to the other while (partially) occupying both lanes simultaneously for some period of time. This corresponds to the same car existing on multiple lanes for some time (studying the actual local curve dynamics is beyond the scope of this paper, but benefits from our modular hierarchical proof structure).

Gradual lane change is modeled by an appearance of i on the new lane ($E(i, l') := 1$) when the lane change starts, then a period of simultaneous existence on both lanes while the car is in the process of moving over, and then, eventually, disappearance from the old lane ($E(i, l) := 0$) when the lane change has been completed and the car occupies no part of the old lane anymore. Consequently, gradual lane change is over-approximated by a series of deletes from all lanes ($\forall l : L \text{ delete}_l^*$) together with a series of appearances on all lanes ($\forall l : L \text{ new}_l^*$). Global highway control with multiple cars moving on multiple lanes and non-deterministic gradual lane changing can be modeled by QHP:

$$\text{ghc} \equiv (\forall l : L \text{ delete}_l^*; \forall l : L \text{ new}_l^*; \forall l : L \text{ ctr}_l^m; \forall l : L \text{ dyn}_l^n)^*$$

Verification. Global highway control ghc is safe, i.e., guarantees collision freedom for multi-lane car control with arbitrarily many lanes, cars, and gradual lane changing.

Theorem 1 (Safety of global highway control ghc). *The global highway control system (ghc) for multi-lane distributed car control is collision-free. This is expressed by the provable formula:*

$$\forall l : L \forall i : C_l(i \ll L_l(i)) \rightarrow [(\forall l : L \text{ delete}_l^*; \forall l : L \text{ new}_l^*; \forall l : L \text{ ctr}_l^m; \forall l : L \text{ dyn}_l^n)^*] \forall l : L \forall i : C_l(i \ll L_l^*(i))$$

For the proof see [23, Appendix A.4]. Note that the constraints on safe lane changing coincide with those identified in Sect. 7 for safe appearance on a lane.

9 Conclusion and Future Work

Distributed car control has been proposed repeatedly as a solution to safety and efficiency problems in ground transportation. Yet, a move to this next generation technology, however promising it may be, is only wise when its reliability has been ensured. Otherwise the cure would be worse than the disease. Distributed car control has been out of scope for previous formal verification techniques. We have presented formal verification results guaranteeing collision freedom in a series of increasingly complex settings, culminating in a safety proof for distributed car control despite an arbitrary and evolving number of cars moving between an arbitrary number of lanes. Our research is an important basis for formally assured car control. The modular proof structure we identify in this paper generalizes to other scenarios, e.g., variations in local car dynamics or changes in system design. Future work includes mechanizing the proof and addressing time synchronization, sensor inaccuracy, curved lanes, and asynchronous sensors.

References

1. Chang, J., Cohen, D., Blincoe, L., Subramanian, R., Lombardo, L.: CICAS-V research on comprehensive costs of intersection crashes. Technical Report 07-0016, NHTSA (2007)
2. Damm, W., Hungar, H., Olderog, E.R.: Verification of cooperating traffic agents. International Journal of Control 79, 395–421 (2006)

3. Dao, T.S., Clark, C.M., Huissoon, J.P.: Distributed platoon assignment and lane selection for traffic flow optimization. In: IEEE IV 2008, pp. 739–744 (2008)
4. Dao, T.S., Clark, C.M., Huissoon, J.P.: Optimized lane assignment using inter-vehicle communication. In: IEEE IV 2007, pp. 1217–1222 (2007)
5. Hall, R., Chin, C., Gadgil, N.: The automated highway system / street interface: Final report. PATH Research Report UCB-ITS-PRR-2003-06, UC Berkeley (2003)
6. Hall, R., Chin, C.: Vehicle sorting for platoon formation: Impacts on highway entry and throughput. PATH Research Report UCB-ITS-PRR-2002-07, UC Berkeley (2002)
7. Hsu, A., Eskafi, F., Sachs, S., Varaiya, P.: Design of platoon maneuver protocols for IVHS. PATH Research Report UCB-ITS-PRR-91-6, UC Berkeley (1991)
8. Ioannou, P.A.: Automated Highway Systems. Springer, Heidelberg (1997)
9. Jula, H., Kosmatopoulos, E.B., Ioannou, P.A.: Collision avoidance analysis for lane changing and merging. PATH Research Report UCB-ITS-PRR-99-13, UC Berkeley (1999)
10. Horowitz, R., Tan, C.W., Sun, X.: An efficient lane change maneuver for platoons of vehicles in an automated highway system. PATH Research Report UCB-ITS-PRR-2004-16, UC Berkeley (2004)
11. Shladover, S.E.: Effects of traffic density on communication requirements for Cooperative Intersection Collision Avoidance Systems (CICAS). PATH Working Paper UCB-ITS-PWP-2005-1, UC Berkeley (2004)
12. Stursberg, O., Fehnker, A., Han, Z., Krogh, B.H.: Verification of a cruise control system using counterexample-guided search. Control Engineering Practice 38, 1269–1278 (2004)
13. Varaiya, P.: Smart cars on smart roads: problems of control. IEEE Trans. Automat. Control 38, 195–207 (1993)
14. Wongpiromsarn, T., Mitra, S., Murray, R.M., Lamperski, A.: Periodically controlled hybrid systems: Verifying a controller for an autonomous vehicle. In: Majumdar, R., Tabuada, P. (eds.) HSCC 2009. LNCS, vol. 5469, pp. 396–410. Springer, Heidelberg (2009)
15. Chee, W., Tomizuka, M.: Vehicle lane change maneuver in automated highway systems. PATH Research Report UCB-ITS-PRR-94-22, UC Berkeley (1994)
16. Johansson, R., Rantzer, A. (eds.): Nonlinear and Hybrid Systems in Automotive Control. Society of Automotive Engineers Inc. (2003)
17. Althoff, M., Althoff, D., Wollherr, D., Buss, M.: Safety verification of autonomous vehicles for coordinated evasive maneuvers. In: IEEE IV 2010, pp. 1078–1083 (2010)
18. Berardi, L., Santis, E., Benedetto, M., Pola, G.: Approximations of maximal controlled safe sets for hybrid systems. In: Johansson, R., Rantzer, A. (eds.) Nonlinear and Hybrid Systems in Automotive Control, pp. 335–350. Springer, Heidelberg (2003)
19. Platzer, A.: Quantified Differential Dynamic Logic for Distributed Hybrid Systems. In: Dawar, A., Veith, H. (eds.) CSL 2010. LNCS, vol. 6247, pp. 469–483. Springer, Heidelberg (2010)
20. Lygeros, J., Lynch, N.: Strings of vehicles: Modeling safety conditions. In: Henzinger, T.A., Sastry, S.S. (eds.) HSCC 1998. LNCS, vol. 1386, pp. 273–288. Springer, Heidelberg (1998)
21. Dolginova, E., Lynch, N.: Safety verification for automated platoon maneuvers: A case study. In: Maler, O. (ed.) HART, pp. 154–170. Springer, Heidelberg (1997)
22. Electronic Proof and Demo, <http://www.l.s.cs.cmu.edu/dccs/>
23. Loos, S.M., Platzer, A., Nistor, L.: Adaptive cruise control: Hybrid, distributed, and now formally verified. Technical Report CMU-CS-11-107, Carnegie Mellon University (2011)
24. Germann, S.: Modellbildung und Modellgestützte Regelung der Fahrzeuflängsdynamik. Fortschrittsberichte VDI, Reihe 12, Nr. 309, VDI Verlag (1997)