Efficiency Analysis of Formally Verified Adaptive Cruise Controllers

Sarah M. Loos[1], David Witmer[2], Peter Steenkiste[3], André Platzer[4]

***Abstract*— We consider an adaptive cruise control system in which control decisions are made based on position and velocity information received from other vehicles via V2V wireless communication. If the vehicles follow each other at a close distance, they have better wireless reception but collisions may occur when a follower car does not receive notice about the decelerations of the leader car fast enough to react before it is too late. If the vehicles are farther apart, they would have a bigger safety margin, but the wireless communication drops out more often, so that the follower car no longer receives what the leader car is doing. In order to guarantee safety, such a system must return control to the driver if it does not receive an update from a nearby vehicle within some timeout period. The value of this timeout parameter encodes a tradeoff between the likelihood that an update is received and the maximum safe acceleration. Combining formal verification techniques for hybrid systems with a wireless communication model, we analyze how the expected efficiency of a provably-safe adaptive cruise control system is affected by the value of this timeout.**

I. Introduction

We present a class of adaptive cruise controllers that safely set the acceleration of the controlled car based on vehicle to vehicle (V2V) communication data from the car ahead. In the design of such a controller, we can take advantage of fast reaction times resulting from V2V communication, allowing cars to drive close together and increasing highway throughputs. However, a wireless transmission may not be received because of interference, physical obstructions, or the distance between the cars being too large. As a result, any controller that depends on V2V communication must also be able to request help from the driver when that communication fails. If the distance between two cars is kept large, then the car controller is more robust to such communication errors because it has more space to maneuver safely, but the probability that a transmission fails increases at larger distances. Furthermore, the throughput of the highway is reduced. For small distances between two cars, communication works more reliably, but there is less room for errors. At close range, as soon as a single wireless message fails to be delivered, the follower car would already have to brake, because slowing down is the only guaranteed safe action at close distance when the car no longer has reliable position and velocity data on the leader car. The same issues arise for systems that ask for driver assistance instead of decelerating automatically, as minimizing driver intervention is one goal of such systems.

We use a symbolic class of adaptive cruise controllers to investigate this tradeoff between efficiency and robustness to communication failures quantitatively. This analysis requires both a safety argument for the hybrid system cruise control and an efficiency argument for a probabilistic communication model. The control model for the physical dynamics of the car is a hybrid system with discrete control decisions and an analysis of their impact on the car's continuous motion. The communication model, instead, depends on the physical relationships like distances, but has a probabilistic nature, because communication attempts may succeed or fail at random according to a corresponding distribution.

In Section IV, we introduce a control function which, given the current position and velocity of the two cars, chooses a new acceleration for the following car. We provide a formal verification proof that our proposed control function does not allow the following vehicle to collide with the leader, under the assumption that if no message update is received within a bounded limit, the driver assumes control of the vehicle. In Section V, we prove that this control function is optimal by showing that any larger choice of acceleration may admit a collision, and would therefore be unsafe. We then use this safe and optimal control function to analyze how changing the time the controller waits before requesting human assistance affects the range of safe acceleration choices. The methodology and results of this analysis are presented in Section VI.

II. Related Work

We would like our adaptive cruise control system to be as efficient as possible, but it must also behave safely, even in the unlikely event that a communication update is not received for a prolonged period. In this paper we work within the context of a controller for which safety has been verified over the full continuous state space which results from all possible discrete actions taken by the controller.

This is a stronger guarantee than what can be given by verification methods that discretize the continuous state space, such as the probabilistic model checker Prism [1]. Other methods for verifying hybrid systems, such as SpaceEx [2], only verify linear hybrid systems, and therefore can not handle the nonlinearities required for setting the maximal

This material is based upon work supported by the National Science Foundation under NSF CAREER Award CNS-1054246, NSF EXPEDITION CNS-0926181, grant no. CNS-0931985, and grant no. CNS-1035800. This research was also supported by the US Department of Transportation's University Transportation Center's TSET grant, award no. DTRT12GUTC11.

[1] S. Loos is with the Computer Science Department, Carnegie Mellon University, Pittsburgh, USA sloos@cs.cmu.edu and is supported by the Department of Energy Computational Science Graduate Fellowship.

[2] D. Witmer is with the Computer Science Department, Carnegie Mellon University, Pittsburgh, USA dwitmer@cs.cmu.edu.

[3] P. Steenkiste is with the Computer Science Department, Carnegie Mellon University, Pittsburgh, USA prs@cs.cmu.edu.

[4] A. Platzer is with the Computer Science Department, Carnegie Mellon University, Pittsburgh, USA aplatzer@cs.cmu.edu.

choice of acceleration as we do in this paper, and which is necessary for analyzing efficiency of the timeout choice.

While the formal verification of adaptive cruise control presented in [3] can and does handle nonlinear hybrid systems, it uses an implicit choice of acceleration. This means that the adaptive cruise control model presented for two cars in [3] is less challenging to prove safe, but its nondeterminisms make it difficult to implement and not well suited for arguing about the efficiency of acceleration choices. The assumptions in [3] also require that the maximum time elapsed between transmission broadcasts be bounded by a known constant. For wireless communication, it is not possible to guarantee that any communication is ever successful, so this assumption is infeasible. Instead, we require the driver to take control of the vehicle anytime the communication delay exceeds a given timeout $\mathcal{T}$. We build on the results presented in [3] to prove safety for a class of controllers that use explicit assignments. We then analyze this class of controllers to discover the optimal timeout $\mathcal{T}$ for passing control of the vehicle back to the driver in case of network failure. This analysis incorporates the probability of successful reception associated with wireless packets sent at varying distances.

Wireless communication between vehicles is a promising tool for improving highway safety. Hartenstein and Laberteaux [4] give an overview of vehicular ad-hoc networks. Jiang et al. [5] propose a specific protocol for safety-related wireless communications between vehicles. The VSC-A report [6] describes an extensive series of experiments testing the performance of wireless V2V communication in a variety of situations that are typically problematic for autonomous vehicle safety systems. Sepulcre and Gozalvez [7] observe that achieving safety in different roadway scenarios could impose very different specifications on the underlying wireless communications system. Meireles et al. [8] experimentally study the effects of physical obstructions, including other vehicles, on packet delivery ratio and signal strength in wireless V2V communication.

Much work has been done on modeling V2V networks. The survey by Stanica, Chaput, and Beylot [9] provides a good overview of current techniques. While modeling wireless networks is relatively well understood, V2V networks pose a unique challenge because vehicles move quickly relative to each other and their environment. The Doppler effect and effects due to the presence of cars, buildings, and other structures could be significant. Dhoutaut et al. [10] propose a simple model that switches between a small number of predefined interference states. Moser et al. [11], on the other hand, give a much more accurate model using raytracing that relies on having a detailed model of the environment and requires much greater computational resources to simulate.

In this paper, we will use the Nakagami fading model [12] to compute the probability that an individual transmission is passed successfully from one car to another as a function of the distance between the two vehicles. We use this model for simplicity. Another more complicated model could be more accurate and our techniques would allow the use of such a model in our efficiency analysis.

III. Preliminaries: Differential Dynamic Logic

Automated car control systems are hybrid systems, which we model by *hybrid programs* (HPs) [13], [14], [15]. HPs are defined by the grammar (α, β are HPs, θ a term, x a variable, and H a formula of first-order logic):

$$\alpha, \beta \ ::= \ x := \theta \mid x_1' = \theta_1, ..., x_i' = \theta_i \,\&\, H \mid x := * \mid ?H \mid \alpha \cup \beta \mid \alpha; \beta \mid \alpha^*$$

The effect of an *assignment* $x := \theta$ is an instantaneous discrete jump assigning θ to x. The effect of *differential equation* $x_1' = \theta_1, ..., x_i' = \theta_i \,\&\, H$ is a continuous evolution where the differential equations $x_i' = \theta_i$ hold *and* (written $\&$ for clarity) formula H holds throughout the evolution (the state remains in the region described by H). Here x' is intended to denote the derivative of the interpretation of the term x over time during continuous evolution. The effect of the random assignment $x := *$ is to non-deterministically pick an arbitrary real number as the value of x.

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* $\alpha; \beta$, HP 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 HPs, we use *differential dynamic logic* $\mathsf{d}\mathcal{L}$ [13], [14], [15] with the grammar:

$$\phi, \psi \ ::= \ \theta_1 = \theta_2 \mid \theta_1 \geq \theta_2 \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \phi \rightarrow \psi \mid [\alpha]\phi \mid \langle\alpha\rangle\phi$$

In addition to all formulas of first-order real arithmetic, $\mathsf{d}\mathcal{L}$ allows formulas of the form $[\alpha]\phi$ with an HP 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 [14], [15] for details.

IV. Verified Adaptive Cruise Control

Hybrid systems have tightly coupled discrete and continuous dynamics. As a result, the consequences of discrete control choices on continuous system dynamics are usually complex and difficult to analyze.

In this paper, we take as a canonical hybrid system an adaptive cruise controller (ACC). The controller uses discrete message updates via V2V about the car ahead to inform discrete acceleration control decisions, which result in continuous changes in position and velocity of the vehicle.

First, we present a similar motivating example as in [3], shown in Figure 1. The leader car l brakes at time $t1$ with its full braking power, -B. But there is a delay before the follower receives a wireless communication update about the behavior of car l, at which point the follow car applies braking power $-b$. But the deceleration is too little and too late, and a collision occurs. The choice of acceleration must be accurate, even when it experiences delays between communication updates from the car ahead.

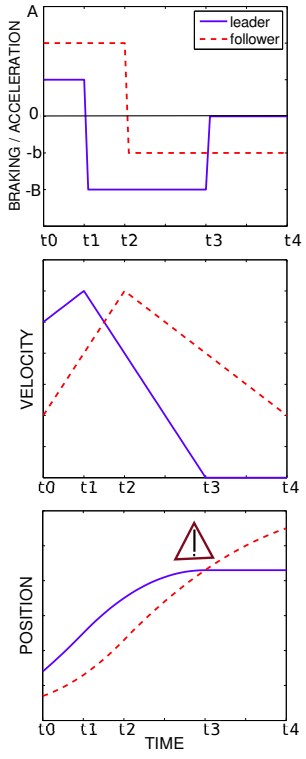

Fig. 1. The wrong choice of acceleration early on may result in an unavoidable collision.

However, the controller cannot drive the car indefinitely while waiting for updates about its environment. If no update is received, it will timeout and request driver assistance. In this system, the cars wirelessly broadcast their position and velocity at a set frequency, but these broadcasts may not be received. In our model of the controller, we define this timeout by the symbolic parameter $\mathcal{T}$. While other parameters in the model are symbolic, many of them are not flexible in a particular implementation of the system, for example the upper and lower bounds on acceleration and braking, which are defined by the physical limits of the car. However, the timeout parameter $\mathcal{T}$ can be set arbitrarily in software. It is crucial to the efficiency of the controller to set $\mathcal{T}$ optimally, however finding the optimal value for $\mathcal{T}$ is nontrivial, as we will discuss in greater detail in Section VI.

In Controller 1, we model two cars driving along a straight road, where the lead car may choose its acceleration arbitrarily (line 4), but the acceleration of the follow car is chosen by an automated control system (line 5). We assume that neither car may travel backward (line 10). The adaptive cruise control system presented in Controller 1 is specified in Differential Dynamic Logic (d$\mathcal{L}$) [13], [14], [15]. The controlled follow car f has state variables x_f, v_f and a_f to represent its position, velocity and acceleration (similarly for the leader car l). The continuous dynamics for f are described by the differential equation system $x_f' = v_f, v_f' = a_f$ (lines 8,9). The position and velocity of the vehicles change continuously, however we don't assume permanent

Controller 1 Verified Adaptive Cruise Control (ACC)

$$(x_f \leq x_l \wedge v_f^2 \leq v_l^2 + 2DB) \rightarrow [\texttt{ACC}](x_f \leq x_l) \tag{1}$$

$$\begin{aligned}
\texttt{ACC} &\equiv (ctrl; dyn)^* && (2)\\
ctrl &\equiv \ell_{ctrl} \;||\; f_{ctrl}; && (3)\\
\ell_{ctrl} &\equiv (a_\ell := *; \;\; ?(-B \leq a_\ell \leq A)) && (4)\\
f_{ctrl} &\equiv a_f := a_f(v_f, v_l, D, \mathcal{T}) && (5)\\
D &\equiv x_l - x_f && (6)\\
dyn &\equiv (t := 0;\; t' = 1, && (7)\\
&\quad x_f' = v_f,\; v_f' = a_f, && (8)\\
&\quad x_\ell' = v_\ell,\; v_\ell' = a_\ell && (9)\\
&\quad \& \; v_f \geq 0 \;\wedge\; v_\ell \geq 0 \;\wedge\; t \leq \mathcal{T}) && (10)
\end{aligned}$$

$$a \equiv \frac{\sqrt{B^2\mathcal{T}^2 - 4Bv_f\mathcal{T} + 8BD + 4v_l^2} - B\mathcal{T} - 2v_f}{2\mathcal{T}} \tag{11}$$

$$b \equiv \frac{-v_f^2}{2(D + \frac{v_l^2}{2B})} \tag{12}$$

$$a_f(v_f, v_l, D, \mathcal{T}) := \begin{cases} A & \text{if } a \geq A \\ 0 & \text{if } v_f = 0 \wedge a \leq 0 \\ a & \text{if } a \geq \frac{-v_f}{\mathcal{T}} \wedge -B \leq a \\ b & \text{if } a < \frac{-v_f}{\mathcal{T}} \wedge -B \leq b \\ -B & \text{o.w.} \end{cases} \tag{13}$$

control over the acceleration, since V2V message updates are only available discretely. In this section, we assume that upon receipt of a packet from the leader car with current values of x_l and v_l, the controller for car f sets its acceleration and maintains it until receiving a subsequent message from l. This behavior is captured by the nondeterministic repetition * in line 2. We assume a maximum acceleration for both cars and denote it by $A > 0$. We also assume a maximum braking power $B > 0$.

In order to automatically control a car, a computer must have specific algorithms for setting acceleration and these algorithms must be guaranteed to keep the two cars separated under all circumstances. The complex interactions between discrete and continuous components make even the simplest control systems challenging to implement safely. The formulas required to calculate an appropriate acceleration in every circumstance quickly become very complex.

Controller 1 presents a formula for setting the acceleration of car f in lines 11-13. It is a function of the relative position D and the velocities v_f, v_l of the two cars. It also relies on the choice of the timeout parameter, $\mathcal{T}$, which determines how long the automated control will wait for an update before requesting assistance from the human driver. In line

11, a is designed to be the greatest acceleration such that if car l is braking maximally, and car f accelerates at rate a for up to $\mathcal{T}$ time and then also brakes maximally, the two cars will not collide. Note that a may be positive or negative. In line 12, b is designed to be the least braking required to bring f to a stop before the point where l would stop, if it is applying maximum braking $-B$.

In line 13, we set the acceleration a_f to be a only if it would not cause the car to stop before time $\mathcal{T}$, otherwise b is chosen. The choice of acceleration is also limited by the physical bounds of the car, A and $-B$. Finally, if the car is stopped, it may either accelerate or remain stopped, but it can not be made to travel backward.

Using the theorem prover KeYmaera [16], we proved the safety property in line 1, which states that while the controller for car f chooses its acceleration based on the formula in line (13), the cars will not collide. To prove this property, we must assume car f is initially behind car l and the state is initially in a controllable region, i.e. if both cars were to immediately brake maximally, they could avoid a collision. This property is expressed by the following formula, which we prove is invariant for `ACC` (i.e. if it holds initially, it continues to hold through all executions of `ACC`):

$$v_f^2 \leq v_\ell^2 + 2DB \tag{14}$$

This assumption can be seen as an antecedent in line 1. This proof of safety requires 334 user interactions, has 661 nodes and takes just 24.2 seconds to prove on a laptop computer. The proof file may be downloaded online from `http://www.ls.cs.cmu.edu/pub/acc/`.

V. Optimality

To prove that the function a_f in Controller 1 is optimal, we show that if a_f were chosen to be $a_f + \varepsilon$ for any $\varepsilon > 0$, there could be a collision. Specifically, we show that in the worst case where l is already braking maximally and car f accelerates at rate $a_f + \varepsilon$ for the maximum duration of $\mathcal{T}$, then even when car f applies maximum braking at time $\mathcal{T}$, it is not able to avoid a collision.

First, we remind the reader of some useful definitions which can be derived in the standard way from kinematic equations and by integrating the ODE in line 8 of Controller 1. Let $x_c[a^t]$ be the position of car c, after accelerating at rate a for time t. Similarly, let $v_c[a^t]$ be the velocity of car c after accelerating at rate a for time t.

$$x_c[(a,t)] = \frac{1}{2}at^2 + v_c t + x_c \qquad v_c[(a,t)] = at + v_c$$

We let $x_c[(b, stop)]$ be the location where car c comes to a stop after decelerating at rate b.

$$x_c[(b, stop)] = x + \frac{v^2}{-2b}$$

From these familiar definitions we can compute the position of car c after it has accelerated at rate a for time $\mathcal{T}$ and then brakes maximally until it comes to a stop:

$$x_c[(a,\mathcal{T});(-B, stop)] = x_{f,a}(\mathcal{T}) + \frac{v_{f,a}(\mathcal{T})^2}{2B} \tag{15}$$

Case 1: $a_f = a = \frac{\sqrt{B^2\mathcal{T}^2 - 4Bv_f\mathcal{T} + 8BD + 4v_l^2} - B\mathcal{T} - 2v_f}{2\mathcal{T}}$

Through algebra and the assumption that $a_f \geq -v_f/\mathcal{T}$ given in line (13) of Controller 1, we can show that $x_{f,a_f}(t_{stop}) = x_{l,-B}(t_{stop})$. So the system is safe in this scenario for acceleration choice a_f. This is not surprising since we derived a formal proof of safety for all scenarios in Section IV. (Note that since we are assuming cars are infinitesimal points, $x_f = x_l$ is still considered safe, but $x_f > x_l$ violates our safety condition from line 1.)

However, when we choose to accelerate at rate $a_f + \varepsilon$, we can derive the following:

$$\begin{aligned}
& x_f[(a_f+\varepsilon,\mathcal{T});(-B, stop)] && (16)\\
= \; & \frac{1}{2}(a_f+\varepsilon)\mathcal{T}^2 + v_f\mathcal{T} + x_f + \frac{((a_f+\varepsilon)\mathcal{T} + v_f)^2}{2B} && (17)\\
> \; & \frac{1}{2}a_f\mathcal{T}^2 + v_f\mathcal{T} + x_f + \frac{(a_f\mathcal{T} + v_f)^2}{2B} && (18)\\
= \; & x_l + v_l^2/(2B) && (19)\\
= \; & x_l[(-B, stop)] && (20)
\end{aligned}$$

The formula in eq. (18) is equal to the formula in eq. (19) by design through the definition of a_f. It can be easily checked by substituting a for a_f. So, we have shown the system is unsafe in this scenario for the acceleration choice $a_f + \varepsilon$ for any $\varepsilon > 0$. As a result, we know that a_f is the maximal acceleration choice available for this case.

Case 2: $a_f = b = \frac{-v_f^2}{2(D + \frac{v_l^2}{2B})}$

The proof of optimality for this case follows similarly to that in Case 1, however it does not depend on $\mathcal{T}$, as the cars come to a stop before $\mathcal{T}$. In Controller 1, $b \leq a$, so in this case $a_f = b \leq a < -v_f/\mathcal{T} \leq 0$. Therefore, from the definition of a_f and the fact that $a_f < 0$, we get

$$x_f + \frac{v_f^2}{-2a_f} = x_l + \frac{v_l^2}{2B}.$$

As a result, there is an epsilon $\varepsilon > 0$ such that,

$$\begin{aligned}
x_f[(a_f+\varepsilon, stop)] &= x_f + \frac{v_f^2}{-2(a_f+\varepsilon)} \\
&> x_l + \frac{v_l^2}{2B} = x_l[(-B, stop)]
\end{aligned}$$

So the system is unsafe in this scenario, and a_f is the maximum acceleration choice available for this case.

All other cases are trivial or extend Cases 1 and 2 trivially.

VI. Efficiency Analysis

Intuitively, we believe that if the follow car can expect to get more frequent communication updates on the position and velocity of the car ahead, it may follow more closely, and therefore improve efficiency. This intuition is quantified by the assignment of a_f in line 11 of Controller 1. When the maximum time between updates, $\mathcal{T}$, is small, the acceleration can be set to a larger value, as demonstrated later in Fig. 3.

Unfortunately, reducing $\mathcal{T}$ is not cost-free, since every time the follow car does not receive a communications update within $\mathcal{T}$, human assistance is required. We want to reduce the frequency of requests for driver intervention, so it might make sense to set the timeout, $\mathcal{T}$, to be large.

To study this tradeoff quantitatively, we will think of the efficiency of the system as the ratio of the control space we can access averaged over the state space. The ratio of the control space that we can access at a single point in the state space is the normalized acceleration $\bar{a}_f$:

$$\bar{a}_f(v_\ell, v_f, D, \mathcal{T}) = \frac{a_f(v_\ell, v_f, D, \mathcal{T}) + B}{A + B}.$$

Since we would like our system to work with little or no required human intervention, we assign the same efficiency value when the system returns control to the driver as we do when it brakes maximally (both will be 0).

Based on our invariant from (14), we define the maximum possible safe velocity of the follower $s_f(v_\ell)$ as

$$s_f(v_\ell) = \min\{(v_\ell^2 + 2DB)^{1/2}, v_{\max}\}.$$

We can now calculate the efficiency Eff_{a_f} of the a_f controller as a function of timeout $\mathcal{T}$ as follows:

$$\text{Eff}_{a_f}(\mathcal{T}) = \frac{1}{Z} \int_0^{D_{\max}} \int_{v_{\min}}^{v_{\max}} \int_{v_{\min}}^{s_f(v_\ell)} \bar{a}_f \; dv_f dv_\ell dD.$$

where $D_{\max}$ is the maximum distance at which the following car can receive updates from the leading car and $v_{\min}$ and $v_{\max}$ are the minimum and maximum permitted velocities. In our computations, we take $v_{\min}$ to be 45 miles per hour and $v_{\max}$ to be 75 miles per hour; this is a typical range for highway driving speeds. We set $D_{\max}$ to be 200 meters. We assume a uniform probability distribution for the initial state over the state space. Z is the volume of the state space, i.e., the integral of 1 over the same region. Bounds on acceleration and braking will be determined by the capabilities of specific vehicles, but in our analysis we use $A = 2$, and $B = 10$.

Note that any choice of acceleration in $[-B, a_f]$ maintains safety. We use $\bar{a}_f$ to measure efficiency because it gives us the ratio of the allowed interval of accelerations that we can safely access at each point in the state space.

However, this is not the whole picture. We expect that we will not be able to receive updates instantaneously, but will instead receive them within some period of time with some probability. The probability that we successfully receive a given transmission will depend on the distance between the vehicles. In their paper [12], Killat and Hartenstein give a model for this reception probability in terms of distance D based on the Nakagami fading model:

$$r(D, \psi) = \left(1 + 3\frac{x^2}{\psi^2} + \frac{9}{2}\frac{x^4}{\psi^4}\right) \exp\left(-3\frac{D^2}{\psi^2}\right),$$

where ψ is the transmission power in meters. For simplicity, we will set $\psi = 100$ for the rest of the paper. We will then let $r(D) = r(D, 100)$. The resulting relationship between distance and reception probability is shown in Fig. 2.

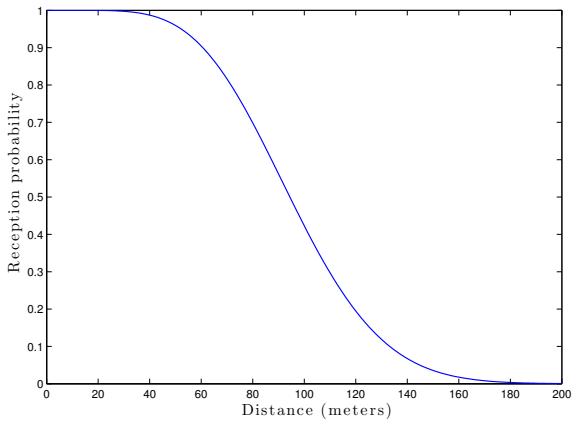

Fig. 2. Reception probability as a function of distance.

For concreteness, we assume that cars broadcast their position and velocity at a frequency of 10 Hz. Our techniques apply for other values, though 10 Hz is used in, e.g., [5], [7]. If the cars stay at a constant distance D, the probability that we receive an update within $\mathcal{T}$ seconds is then

$$1 - (1 - r(D))^{\lfloor 10\mathcal{T} \rfloor}.$$

This formula is not applicable, however, because the distance does not stay constant when both cars move with different velocities or different accelerations. To account for this, we use the initial position, velocity, and acceleration of each car to recalculate the distance between the cars each time an update is sent (recall that $x_c[(a, t)]$ is the position of car c after accelerating at rate a for time t, and that position, velocity and acceleration appear in these terms). So, the probability that we receive an update within $\mathcal{T}$ seconds is

$$p(\mathcal{T}, D, a_\ell, a_f, v_\ell, v_f) = 1 - \prod_{i=1}^{\lfloor 10\mathcal{T} \rfloor} \left(1 - r\left(x_l[(a_l, \frac{i}{10})] - x_f[(a_f, \frac{i}{10})]\right)\right).$$

We take the acceleration of the follow car to be the value a_f returned by the a_f controller. The acceleration of the lead car is not known, so we instead take the average reception probability over all choices of acceleration for the lead car (the formula is easily modified to assume constant velocity, or any probability distribution over acceleration choices):

$$\bar{p}(\mathcal{T}, D, a_f, v_\ell, v_f) = \frac{1}{A + B} \int_{-B}^{A} p(\mathcal{T}, D, a_\ell, a_f, v_\ell, v_f) \; da_\ell.$$

Now we can define the quantity $\text{Eff}_{\text{rec}}(\mathcal{T})$, which captures the average likelihood over the state space that we receive an update within timeout $\mathcal{T}$, as follows:

$$\text{Eff}_{\text{rec}}(\mathcal{T}) = \frac{1}{Z} \int_0^{D_{\max}} \int_{v_{\min}}^{v_{\max}} \int_{v_{\min}}^{s_f(v_\ell)} \bar{p} \; dv_f dv_\ell dD.$$

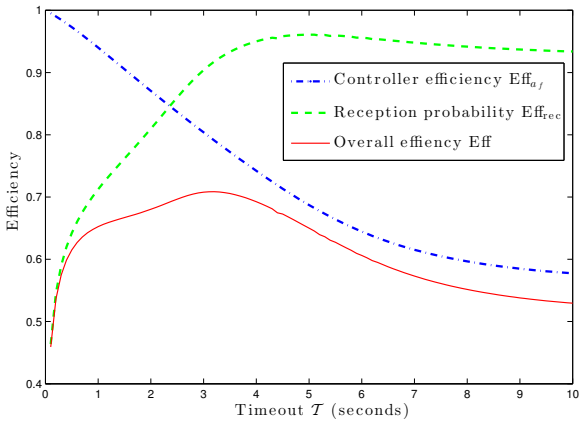

Fig. 3. Reception probability Eff_{rec}, controller efficiency Eff_{a_f}, and overall efficiency Eff as a function of timeout $\mathcal{T}$.

Combining the reception probability and the acceleration, we can calculate the expected efficiency of the whole system:

$$\text{Eff}(\mathcal{T}) = \frac{1}{Z} \int_0^{D_{\max}} \int_{v_{\min}}^{v_{\max}} \int_{v_{\min}}^{s_f(v_\ell)} \bar{a}_f \bar{p} \; dv_f dv_\ell dD.$$

We show the three functions Eff_{a_f}, Eff_{rec}, and Eff in Fig. 3. Eff_{a_f} decreases as $\mathcal{T}$ decreases, since a longer timeout forces the controller to make more conservative decisions. Eff_{rec} initially increases with $\mathcal{T}$, as the probability we successfully receive an update increases for a longer timeout. At higher values of $\mathcal{T}$, the following car may not receive an update from the leading car for a longer period of time. Even though we do not know the behavior of the leading car, the controller must take into account the possibility that the leading car has been continuously applying maximum braking since the last update was received. To ensure safety, the following car must then choose lower acceleration values. However, in much of the state space, the leading car is not actually braking, meaning that the distance to the leading car is increasing in much of the state space. This causes the reception probability to decrease slightly. Eff starts low because we are less likely to receive an update in a short amount of time, then increases because we are more likely to receive updates, and then decreases because the longer timeout forces the controller to make more conservative decisions. It achieves a maximum value of 0.709 at a timeout $\mathcal{T} = 3.2$ seconds.

VII. Conclusions and Future Work

In this paper, we present a symbolic controller for automated car control on a straight road. We identify its safe region and formally verify that it prevents collisions. This strong formal guarantee is needed to ensure the safety-critical functioning of the system. We then investigate a particular instance of the controller with exact values for parameters such as update frequency, signal strength and maximum braking and acceleration. We find the timeout value for communication updates that maximizes the range of safe accelerations over the state space. Although we stepped through the analysis using a relatively simple model of wireless communication, our method is general enough that it could be applied using a more complex communication model tailored to the system being analyzed.

Since the probability of receiving a message from the car ahead depends heavily on the distance between the two cars, another reasonable controller might be one which adjusts the timeout as a function of the distance between the two vehicles. This could be interesting future work, and due to the generality of our analysis techniques, the primary challenge will be in the verification step. We may also be able to allow a closer following distance if we incorporate more realistic probability distributions over the state space, rather than uniform distribution. Future work could also incorporate more cars on the road, increasing the difficulty of the formal proof of safety and necessitating a more complex model of the communication network to represent the large number of collisions of broadcast packets.

References

[1] Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In Gopalakrishnan, G., Qadeer, S., eds.: CAV. Volume 6806 of LNCS., Springer (2011) 585–591

[2] Frehse, G., Guernic, C.L., Donz, A., Cotton, S., Ray, R., Lebeltel, O., Ripado, R., Girard, A., Dang, T., Maler, O.: Spaceex: Scalable verification of hybrid systems. In Gopalakrishnan, G., Qadeer, S., eds.: CAV. (2011) 379–395

[3] Loos, S.M., Platzer, A., Nistor, L.: Adaptive cruise control: Hybrid, distributed, and now formally verified. In Butler, M., Schulte, W., eds.: FM. LNCS, Springer (2011)

[4] Hartenstein, H., Laberteaux, K.: A tutorial survey on vehicular ad hoc networks. Communications Magazine, IEEE **46**(6) (2008) 164–171

[5] Jiang, D., Taliwal, V., Meier, A., Holfelder, W., Herrtwich, R.: Design of 5.9 GHz DSRC-based vehicular safety communication. Wireless Communications, IEEE **13**(5) (2006) 36–43

[6] Ahmed-Zaid, F., Bai, F., Bai, S., Basnayake, C., Bellur, B., Brovold, S., Brown, G., Caminiti, L., Cunningham, D., Elzein, H., et al.: Vehicle safety communications–applications (VSC-A) final report: Appendix volume 1 system design and objective test. Technical report (2011)

[7] Sepulcre, M., Gozalvez, J.: On the importance of application requirements in cooperative vehicular communications. In: Wireless On-Demand Network Systems and Services (WONS). (2011) 124–131

[8] Meireles, R., Boban, M., Steenkiste, P., Tonguz, O., Barros, J.: Experimental study on the impact of vehicular obstructions in VANETs. In: Vehicular Networking Conference (VNC), IEEE. (2010) 338–345

[9] Stanica, R., Chaput, E., Beylot, A.L.: Simulation of vehicular ad-hoc networks: Challenges, review of tools and recommendations. Computer Networks **55**(14) (2011) 3179 – 3188

[10] Dhoutaut, D., Régis, A., Spies, F.: Impact of radio propagation models in vehicular ad hoc networks simulations. In: Proceedings of the 3rd international workshop on Vehicular ad hoc networks. VANET '06, New York, NY, USA, ACM (2006) 40–49

[11] Moser, S., Kargl, F., Keller, A.: Interactive realistic simulation of wireless networks. In: Interactive Ray Tracing, 2007. RT '07. IEEE Symposium on. (2007) 161 –166

[12] Killat, M., Hartenstein, H.: An empirical model for probability of packet reception in vehicular ad hoc networks. EURASIP Journal on Wireless Communications and Networking **2009**(1) (2009) 721301

[13] Platzer, A.: Differential dynamic logic for hybrid systems. J. Autom. Reas. **41**(2) (2008) 143–189

[14] Platzer, A.: Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics. Springer, Heidelberg (2010)

[15] Platzer, A.: Logics of dynamical systems. In: LICS, IEEE (2012) 13–24

[16] Platzer, A., Quesel, J.D.: KeYmaera: A hybrid theorem prover for hybrid systems. In Armando, A., Baumgartner, P., Dowek, G., eds.: IJCAR. Volume 5195 of LNCS., Springer (2008) 171–178