Safe and efficient collision avoidance control for autonomous vehicles

Qiang Wang, Dachuan Li
Research Institute of Trustworthy Autonomous Systems
Southern University of Science and Technology
Shenzhen, China

Joseph Sifakis
Verimag
Université Grenoble Alpes
Grenoble, France

***Abstract*—We study a novel principle for safe and efficient collision avoidance that adopts a mathematically elegant and general framework making as much as possible abstraction of the controlled vehicle's dynamics and of its environment. Vehicle dynamics is characterized by pre-computed functions for accelerating and braking to a given speed. Environment is modeled by a function of time giving the free distance ahead of the controlled vehicle under the assumption that the obstacles are either fixed or are moving in the same direction. The main result is a control policy enforcing the vehicle's speed so as to avoid collision and efficiently use the free distance ahead, provided some initial safety condition holds.**

The studied principle is applied to the design of a synchronous controller. We show that the controller is safe by construction. Furthermore, we show that the efficiency strictly increases for decreasing granularity of discretization. We present the implementation and experimental evaluations in the Carla autonomous driving simulator and investigate various performance issues.

***Index Terms*—Safe and efficient collision avoidance, Autonomous vehicles, Model based design**

I. Introduction

As a fundamental requirement for autonomous vehicle control, the problem of collision avoidance has been widely investigated using a variety of approaches and frameworks. The assumptions underlying the adopted frameworks vary regarding the level of modeling of the dynamics of the controlled vehicle, the number of vehicles and the type of their trajectories or the nature of the controller stimuli. Control-based techniques [1], [7], [8] typically focus on the collision avoidance for adaptive cruise control. Such techniques allow achieving optimality without however providing safety guarantees. Another line of work applies formal methods and model-based design, e.g., reachability analysis [5], [9], Responsibility-Sensitive Safety (RSS) model [11], logic-based controller synthesis [6], [10], as well as the design of specific safety supervision mechanisms [3], [4]. Although such work can guarantee correctness by construction, they may result in policies that achieve strict safety at the expense of efficiency or performance.

We propose a novel principle for safe and efficient collision avoidance. We adopt a mathematically simple and general framework making abstraction of the controlled vehicle's specific dynamics and of its environment, and using only three functions: (1) the free distance function $F(t)$ which determines for the vehicle the estimated free distance from the closest obstacle ahead at time t; (2) the accelerating function $A(V, v)$ which gives the distance travelled by the vehicle when accelerating from initial speed V to speed v; (3) the braking function $B(V, v)$ which gives the distance travelled by the vehicle when braking from V to speed v ($v < V$). The principle consists in the application of a simple induction rule. If at some time t the speed of the vehicle with respect to the distance $F(t)$ is safe, i.e. $B(V, 0) \leq F(t)$, then the speed will be controlled to remain safe under the assumption that $F(t)$ does not change faster than the vehicle can brake. This assumption always holds when the obstacles ahead are fixed or move in the same direction as the controlled vehicle. Furthermore, if safety can be guaranteed for speed V and $B(V, 0) \leq F(t)$ then in order to efficiently use the available space $F(t) - B(V, 0)$, we apply an accelerating/braking policy: we accelerate to a certain speed $v > V$, from which it is still possible to safely brake. So efficiency boils down to computing the maximum target speed v such that $0 \leq F - (A(V, v) + B(v, 0))$. The control principle consists in the dynamic application of the accelerating/braking policy for a set of possible speed levels between speed 0 and the limit speed of the vehicle.

We provide a synchronous controller for safe and efficient collision avoidance, which is driven by periodically sampled values of the free distance F. We prove that this controller is safe and efficient. We also present the implementation and experimental evaluations in the Carla autonomous driving simulator and investigate various performance issues.

Our approach is characterized by the following:

1) It makes abstraction of the vehicle dynamics through the use of accelerating and braking functions that provide all the information needed for safe and efficient control. These functions are a kind of contract between the controller and the controlled vehicle. Their use frees us from the obligation to model vehicle dynamics. Furthermore, it leaves completely open the way features related to comfort such as the jerk profile are implemented.
2) Although we consider the problem in one dimension and the environment is modeled by a free distance function $F(t)$, the result can be easily extended to two dimensions. In that case $F(t)$ and $B(v, V)$ become areas and the safety test consists in checking their inclusion.
3) The control principle is robust and easy to adapt to varying uncertainty in the measurement of F or in the

978-1-7281-9148-5/20/$31.00 © 2020 IEEE

estimation of the functions A and B.

4) The proposed implementation does not have any specific hardware requirements and require very limited computing resources as they combine pre-computed control policies.
5) Finally, the adopted control principle is simple and inductive: if at some step the distance is safe then a speed increase by some quantity will not jeopardize safety. This induction hypothesis is used to prove correctness.

The paper is a short version of [12] where proofs and more details experimental results are provided. Section II presents the framework and the principle of safe and efficient collision avoidance control. Section III presents the design of the collision avoidance controller. Section IV presents the implementation and performance evaluations using the Carla simulator. Section V concludes about the relevance of the results and outlines directions for future work.

II. Safe and efficient collision avoidance control

The aim is to control the movement of a vehicle travelling in a one-way lane, so as to 1) avoid collision with other objects that may be fixed or moving in the same direction (i.e., safety); and 2) use the available free distance ahead in the best possible manner to minimize travelling time (i.e., efficiency).

Our work relies on a mathematically abstract framework characterized by three functions. We denote by v the speed variable of the vehicle and by V its initial speed.

- The function $F(t)$ gives the free distance at time t between the controlled vehicle and the closest obstacle ahead, which either moves in the same direction or is stopped.
- The braking function $B(V,v)$ is a partial function defined in the interval $0 \leq v \leq V$. It gives the distance travelled by the controlled vehicle when braking from the initial speed V to a target speed v. In Fig.1 it is graphically illustrated by the green curves. When the target speed $v = 0$ (i.e, the vehicle brakes to stop), this function is abbreviated as $B(V)$ for simplicity.
- The accelerating function $A(V,v)$ is a partial function defined in the interval $V \leq v \leq V_L$, where V_L is a given limit speed for each vehicle. It gives the distance travelled by the vehicle when accelerating from an initial speed V to a target speed v. In Fig.1 it is graphically illustrated by the black curves.

We make no specific assumptions about the implementation of accelerating and braking functions, e.g. whether acceleration and deceleration are constant or variable. Nonetheless, we require that the following additivity and strict monotonicity properties hold.

- $B(V,V) = 0$ and $A(V,V) = 0$.
- Additivity property:

$$B(V,v_1) + B(V_1,v_2) = B(V,v_2), \quad where \quad v_1 = V_1$$

$$A(V,v_1) + A(V_1,v_2) = A(V,v_2), \quad where \quad v_1 = V_1$$

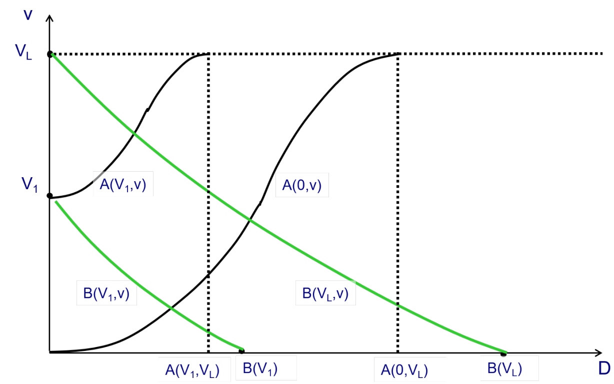

Fig. 1. Braking and acceleration distance functions (D is the distance travelled and v is the speed)

- Strict monotonicity:

$$B(V,v_1) < B(V,v_2), \quad when \quad v_1 < v_2$$

$$A(V,v_1) < A(V,v_2), \quad when \quad v_1 < v_2$$

We progressively study the safe and efficient collision avoidance problem for a vehicle moving in a one-way lane. We first study the problem for a stationary obstacle ahead. Then we study algorithms that solve the problem for dynamically changing free distance. We assume that the movement is controlled using commands for accelerating and braking from a speed V to some target speed v whose effect is modeled by the functions $A(V,v)$ and $B(V,v)$, respectively.

A. Control for safety

The basic idea for avoiding collision is to moderate the speed of the vehicle and anticipate the changes of the free space ahead so as to have enough distance and time to adjust and brake. If the vehicle moves with speed V at time t, then for safety the free space ahead $F(t)$ should be longer than the braking distance $B(V)$, which is the minimal safe braking distance for speed V. The Theorem below formalizes this idea.

Theorem 1. *If at time t the speed V_t of the vehicle is safe with respect to $F(t)$, i.e., $B(V_t) \leq F(t)$ and for any time $t + \triangle t$ it is possible to set the speed to a value $V_{t+\triangle t}$ such that the condition $F(t) - F(t+\triangle t) \leq B(V_t) - B(V_{t+\triangle t})$ holds, then the vehicle is always safe.*

This theorem suggests a simple and safe control policy that ensures collision freedom. For any time t, the vehicle only needs to keep track of the free distance ahead $F(t)$ and check in real-time whether $F(t)$ is greater than the minimal safe braking distance $B(V_t)$ for the current speed V_t. It starts braking as soon as $F(t)$ reaches the minimal safe braking distance. In this way, it is guaranteed that if the obstacles ahead do not move in the opposite direction, no collision would happen.

B. Achieving efficiency for fixed obstacles

The above result provides a basis for ensuring collision freedom. Nonetheless, it leaves open the question of how the vehicle can efficiently use the available distance ahead by minimizing the travelling time. What would be an efficient driving policy when the free headway distance is greater than the minimal safe braking distance? We consider that a policy defines the speed function $v(t)$ in response to a free distance $F(t)$. An Accelerating/Braking policy (A/B policy) is a policy of accelerating first to some speed and then braking. Similarly, an Braking/Accelerating policy (B/A policy) is the policy of braking first to some speed and then accelerating. A Constant speed/Braking policy (C/B policy) is the policy of moving at constant speed and then braking. A policy is safe if the relative distance between the controlled vehicle and the obstacle ahead is positive. It is efficient if increasing the speed value $v(t)$ enforced by the policy at any point would compromise safety.

The problem is to minimize the travelling time for a given distance, which implies to maximize the average speed. Consider the scenario where the speed of the vehicle is V and there is a stationary obstacle ahead at distance F, which is greater than the braking distance $B(V)$. The application of an A/B policy consists in computing an appropriate target speed $v, V < v \leq V_L$, accelerate the vehicle to v and then brake to full stop. To ensure collision freedom, the total travelled distance $A(V, v) + B(v)$ must be such that $A(V, v) + B(v) \leq F$. The maximal target speed is given by the condition $v_M = \max\{v \mid F \geq A(V, v) + B(v)\}$. Such a speed exists as both acceleration and braking functions are monotonically increasing with respect to the target speed v. Notice that either $v_M \leq V_L$ and $F = A(V, v_M) + B(v_M)$ or $v_M = V_L$ and $F > A(V, v_M) + B(v_M)$.

As an example, for motion at constant acceleration and deceleration (a and b, respectively), we have $A(V, v) = v * (v - V)/a + (v - V)^2/2 * a$ and $B(v) = v^2/2 * b$. Then the safety condition becomes $F \geq v * (v - V)/a + (v - V)^2/2 * a + v^2/2 * b$, from which we deduce $v \leq \sqrt{(2 * a * b * F + b * V^2)/(a + b)}$. Thus the maximal target speed $v_M = \sqrt{(2 * a * b * F + b * V^2)/(a + b)}$. As we require that $v \geq V$, we have $F \geq V^2/2 * b = B(V)$ and thus the maximal target speed always exists. Let v_F denote the speed reached by accelerating along distance $F(t)$, i.e., $v_F^2 - V^2 = 2 * F(t) * a$, then the formula can be simplified as $v_M = v_F * \sqrt{b/(a + b)}$.

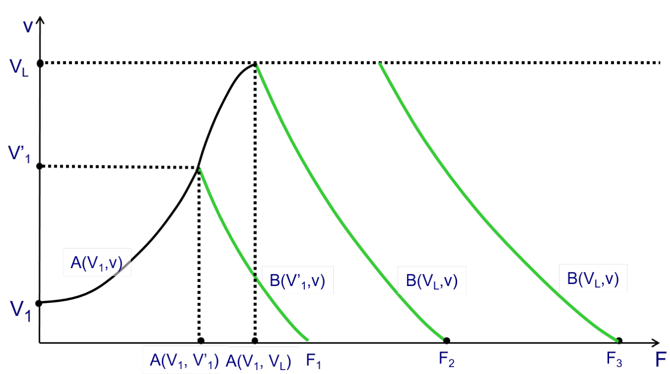

Fig. 2. The A/B control policy for different values of the free distance F ahead

Fig. 2 illustrates the A/B control policy where F is the free distance ahead and v is the speed of the controlled vehicle. The green curves illustrate braking phases and the black the accelerating phase from an initial speed V_1. For $F = F_1$, the maximal target speed V_1' is less than the limit speed V_L. The A/B policy consists in accelerating to V_1', and then braking until the vehicle stops having travelled exactly distance F_1. If the free distance ahead is $F = F_2$, the maximal target speed will be the limit speed V_L. The A/B policy will similarly accelerate first to the limit speed and then brake to stop at F_2. Finally, if $F = F_3 > F_2$, then after accelerating to the limit speed V_L, the vehicle will maintain constant speed V_L for distance $F_3 - A(V_1, V_L) - B(V_L)$ and then brake for the remaining distance to stop at F_3.

Theorem 2. *If the speed V of the vehicle is safe with respect to F, i.e.,$B(V) \leq F$, then the A/B policy is always safe and efficient for F.*

The above result implies that for the given free distance F, the A/B policy is the most efficient and that from the given initial speed there is a maximal speed that minimizes the travel time of F.

III. Controller design for collision avoidance

A. The control principle

We study a control principle for collision avoidance based on the above results. We consider that the vehicle speed can change between a finite set of increasing levels $v_0, v_1, ..., v_n$, where n is a constant, $v_0 = 0$ and v_n equals to the limit speed v_L. The triggering of acceleration and braking from one level to another is controlled according to the free distance ahead and based on bounds computed as follows, for each speed level $v_i, i \in [1, n]$,

- $B_i = B(v_i)$ is the minimal safe braking distance needed for the vehicle to fully stop from speed v_i;
- $D_i = A(v_{i-1}, v_i) + B(v_i)$ is the minimal safe distance needed for the vehicle to apply an A/B policy accelerating from speed v_{i-1} to v_i and then braking from v_i to stop.

We show that the following function specifies the highest safe speed level v as a function of the current speed of the vehicle V and the free space ahead F, provided that their initial values V_0 and F_0 are such that $B(V_0) \leq F_0$.

$$v = Control(F, V)$$

$$v = \begin{cases} v_{i+1} & \text{when} \quad V = v_i \wedge F = D_{i+1} \\ v_{i-1} & \text{when} \quad V = v_i \wedge F = B_i \\ v_i & \text{when} \quad V = v_i \wedge D_{i+1} > F > B_i \end{cases}$$

Note that this control principle is purely functional. It assumes that changes of the free distance ahead F can be continuously monitored to instantaneously produce correspomding speed changes.

Fig.3 illustrates the principle for $n = 4$ speed levels. As the value of F increases, the speed of the vehicle switches between levels. Safety is preserved by construction. The vehicle can

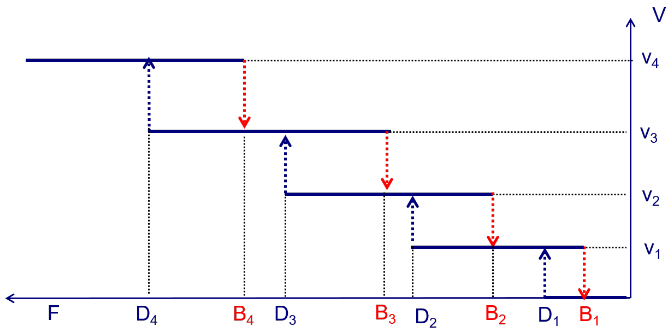

Fig. 3. Illustration of the collision avoidance principle for $n = 4$

accelerate to a higher level, if it can safely and efficiently use the available distance by applying an A/B policy. It brakes to a lower level if the available distance reaches the bound for safe braking.

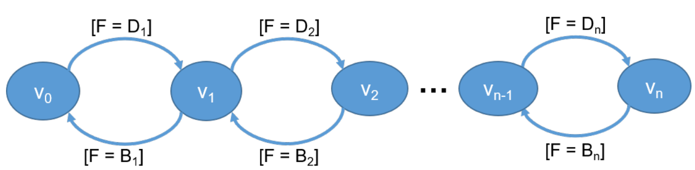

Fig. 4. Automaton modelling the collision avoidance principle

Fig.4 provides a scheme for the computation of $Control(F, V)$ in the form of a finite state automaton. The locations correspond to traveling at constant speeds $v_0, ..., v_n$. The transitions model instantaneous acceleration and braking steps triggered by conditions involving the free distance F and the precomputed bounds B_i and D_i. If the control location is v_i and the free distance ahead equals to the minimal safe acceleration distance (i.e., $F = D_{i+1}$), then the automaton moves to location v_{i+1} after the speed is accelerated to v_{i+1}. If the free distance ahead reaches the minimal safe braking distance (i.e., $F = B_i$), then the automaton moves to location v_{i-1} after the speed is decelerated to v_{i-1}. Recall that $B_i = B_{i-1} + B(v_i, v_{i-1})$. Thus, after braking to v_{i-1} there is still enough space for safe braking. Note that checking point conditions makes sense because F has no jumps and computation is instantaneous. If none of the triggering conditions holds, then the free distance ahead F is such that $B_i < F < D_{i+1}$. The automaton stays at location v_i and the speed remains unchanged.

Note that the automaton of Fig.4 cannot be implemented as a controller because we assume that F is continuously observable and changes of the controlled speed are instantaneous. In the next section, we show how to design a practical controller by refining this automaton.

Theorem 3. *The collision avoidance principle is safe. Moreover, its efficiency is strictly increasing for increasing number of speed levels n.*

As explained, computing the exact value of the optimal speed for a given distance may be costly. Considering discrete speed levels allows pre-computing for each level both the minimal safe braking distance and the minimal safe accelerating distance between levels. In that manner, we avoid the computational complexity of adjusting in real time the vehicle speed.

B. Synchronous controller design

We propose a synchronus controller applying the presented collision avoidance principle. It is driven by periodic updates of the free distance variable F for an adequately chosen period. We have also studied an asynchronous controller where the free distance variable F is updated sporadically [12].

The controller interacts with its controlled environment (the vehicle) through input and output events. The output s is a state variable indicating the currently applied command (i.e., accelerating, braking or constant speed). The input event $UpdateF$ receives the periodic measurement F' of the free distance with period T, while input events ca and cb signal the completion of the accelerating and braking command respectively. Initially, the speed v of the vehicle is set to a level v_i that is safe with respect to the initial distance F (i.e., $F \geq B(v_i)$).

The controller is a refinement of the ideal controller where we assumed that speed changes were instantaneous. Its is described by the following set of guarded commands and also depicted as an extended automaton for the sake of clarity in Fig.5.

$$
\begin{array}{l}
do \\
\Box\ \exists i \in [1, n].s = Csp(v_i) \wedge F' \geq D'_{i+1} \\
\qquad \rightarrow s := Ac(v_i, v_{i+1});\ ba \\
\Box\ \exists i \in [1, n].s = Csp(v_i) \wedge B''_i \geq F' \geq B'_i \\
\qquad \rightarrow s := Br(v_i, v_{i-1});\ bb \\
\Box\ \exists i \in [1, n].ca \wedge s = Ac(v_i, v_{i+1}) \rightarrow s := Csp(v_{i+1}) \\
\Box\ \exists i \in [1, n].cb \wedge s = Br(v_i, v_{i-1}) \rightarrow s := Csp(v_{i-1}) \\
\Box\ UpdateF \rightarrow F' := F \\
od
\end{array}
$$

For guarded commands we adopt the usual semantics: whenever the condition on the left hand side holds, the actions on the right hand side are executed. Note that the input events appear as conditions while the output event appear as actions. The variable s keeps track of the kinematic state of the vehicle that is abstracted by the control states **Ac** (accelerating), **Br** (braking) and **Csp** (moving with constant speed).

We denote by $Ac(v_i, v_{i+1})$, $Br(v_i, v_{i-1})$ and $Csp(v_i)$ the commands of accelerating from speed v_i to v_{i+1}, braking from v_i to v_{i-1} and moving with speed v_i, respectively. When the vehicle is moving with constant speed, transition **UpdateF** is triggered periodically to receive the most recent measurement of F. Once the triggering condition of accelerating (braking) is met, transition **ba** (**bb**) is taken to initiate the command and move to location **Ac** (**Br**) waiting for its completion.

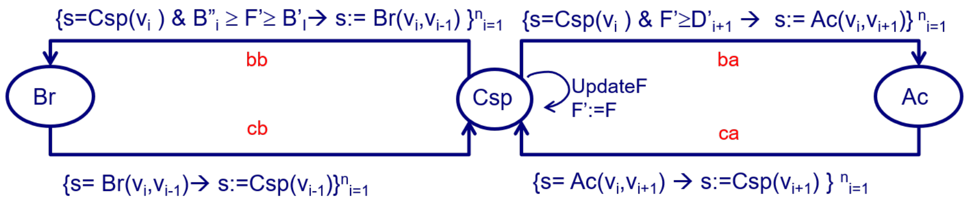

Fig. 5. Extended automaton modelling the synchronous controller

We do not make any assumption about the time spent at locations **Ac** and **Br**. We simply assume that the distances needed for accelerating and braking are $A(v_{i-1}, v_i)$ and $B(v_i, v_{i-1})$, respectively. We explain below how the guards of the controllable transtions bb and ba are computed.

We estimate the maximal safe approximations of the triggering conditions $F \geq D_i$ and $F = B_i$ of the ideal controller in terms of F', the most recently updated value of F. When the vehicle moves at speed v_i, the variables F and F' satisfy a relation of the form $F = F' - k_i(t)$, where $k_i(t) = v_i * (t \ mod \ T)$. That is $k_i(t) = 0$ when F is updated and $k_i(t) < v_i * T$. We assume that T is small enough so that $D_i - v_i * T \geq B_i$, that is, we do not miss the braking threshold value B_i in a period. This is reasonable given that in practice the updating period T is usually less than 50 milliseconds. Notice that the minimal value of F will be reached for $F = F' - v_n * T$. Thus, it is enough to require that $F' \geq D_i + v_n * T$ holds for accelerating and that $B_i + 2 * v_n * T \geq F' \geq B_i + v_n * T$ holds for braking. So we adjust the triggering bound for accelerating to $D'_i = D_i + v_n * T$ and the least and upper bounds of the interval triggering a braking to $B'_i = B_i + v_n * T, B''_i = B_i + 2 * v_n * T$.

Theorem 4. *The synchronous controller yields safe control policies for collision avoidance.*

As for the previous theorem, we can prove that the efficiency of the controller strictly increases for increasing number of speed levels n. Furthermore, the efficiency depends on how frequently F is updated as the accelerating and braking conditions take into account the uncertainty about the values of F. Thus, the controller may not be able to fully utilize actually available free distance. As explained for synchronous controller, the loss in efficiency depends on the value of $v_n * T$.

IV. Experimental Evaluations

We have implemented the synchronous controller in the open-source autonomous driving simulator Carla [2]. In the experiments, we consider scenarios where the controlled vehicle is driving towards a moving vehicle ahead. The speed of the front vehicle is described by the periodic function $v_f(t) = v_{f0} + v_{f0} * sin(\omega * t)$, where $\omega = 2 * \pi / T_f$, T_f is the period of this speed function and v_{f0} is a constant. We choose $v_{f0} = 14 \ m/s$, and thus the speed of the front vehicle changes in the interval $[0, 28 \ m/s]$ (i.e., $[0, 100.8 \ km/h]$). We set the limit speed of the controlled vehicle to be $32 \ m/s$ (i.e., $115.2 \ km/h$). The initial distance between the two vehicles is $F(0) = 5m$ and the initial speed of the controlled vehicle is 0. Thus, the controlled vehicle is initially at a safe state. The accelerating and braking rates of the two vehicles are both constant $a = b = 2 \ m/s^2$.

In order to evaluate the performance and the quality of the controller, we measure both the speed changes of the controlled vehicle and the relative distance between the two vehicles, reflecting the occupancy of the road. The smaller the distance is, the higher the road occupancy is. We consider that the free distance is the relative distance increased by the braking distance of the front vehicle. We perform the evaluations with respect to the period T_f of v_f. For experimental purposes, the safe accelerating and braking distances for eight speed levels $v[8] = \{4, 8, 12, 16, 20, 24, 28, 32\}$ are pre-computed for constant accelerating and braking rates $a = b = 2 \ m/s^2$.

First, we evaluate how T_f affects the performance of the controller for two different values $T_f = 10 \ s$ and $T_f = 30 \ s$. We assume that the environment updates the free distance variable with period $T = 0.02 \ s$. We take into account in the evaluation of the free distance the braking distance of the front vehicle travelling with speed $v_f(t) = v_{f0} + v_{f0} * sin(\omega * t)$. At time t its braking distance is $v_f(t)^2 / 2 * b_f$ for a constant braking rate b_f. We take $b_f = 5m/s^2$ for experimental purposes. Then the corresponding safe accelerating and braking distances are $D'_i = D_i - v_f(t)^2 / 2 * b_f$ and $B'_i = B_i - v_f(t)^2 / 2 * b_f$.

In Fig.6 we compare the results for $T_f \in \{10 \ s, 30 \ s\}$ with sensing period $T = 0.02 \ s$ and speed level $n = 8$. The top two figures compare the dynamics of the relative distance, which is periodic with period T_f in the steady regime. It decreases for increasing period T_f. In fact, for slower speed changes, the controller has more time to adjust the movement of the controlled vehicle and can better utilize the available distance. We can also see that when taking into account the braking distance of the front vehicle, the relative distance between the two vehicles becomes much smaller. For instance, the minimal relative distance decreases from $20.11 \ m$ for $T_f = 30 \ s$ to $11.26 \ m$. The performance improvement can also be observed from the speed diagrams shown in the bottom of Fig.6. Finally, the speed range of the controlled vehicle becomes larger and the maximal speed increases from $16 \ m/s$ to $20 \ m/s$ for $T_f = 10 \ s$.

V. Conclusions and Future Work

The paper presents a novel framework and approach for safe and efficient collision avoidance for self-driving vehicles. The framework is model-based and assumes that control policies

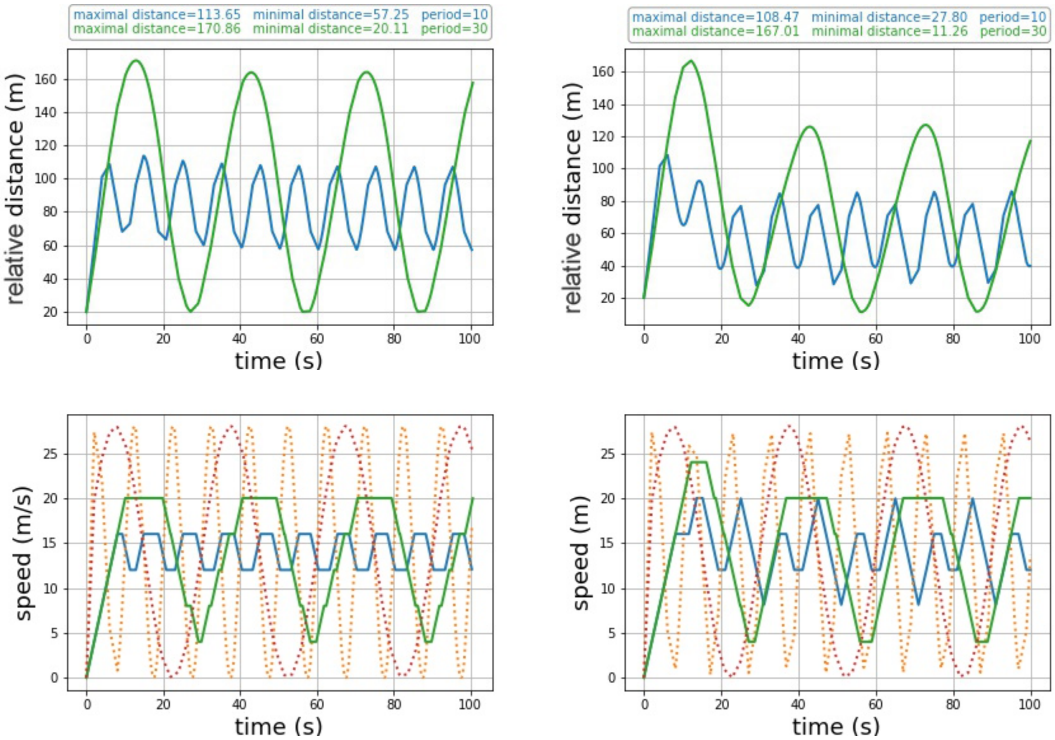

Fig. 6. Simulation results for the synchronous controller with and without considering the braking distance of the front vehicle (right and left part, respectively) for $T_f \in \{10\ s, 30\ s\}$.

are implemented as the application of acceleration, braking and constant speed commands. The presented algorithms do not make any assumption about the dynamics of the controlled vehicle except that there are functions giving the traveled distance when the speed of the vehicle changes by some quantity. Additionally, the assumptions about the vehicle's environment are minimal as it is described by a function $F(t)$ giving at any time the free available distance ahead.

The assumption that all vehicles move in the same direction as the controlled vehicle does not limit the generality of our approach. The same algorithm can also be adapted to two-dimensional movement. In that case, the function $F(t)$ can be defined as the maximal convex area containing the controlled vehicle and such that all the obstacles are outside this area. The distances $A(V, v)$ and $B(V, v)$ for initial and target speeds respectively are also replaced by adequately approximated convex areas so that the safety test boils down to area inclusion that can be efficiently decided.

This work is part of a project on the design of safe and efficient autopilots for autonomous vehicles. Future developments include the adaptation of this algorithm to two-dimension movements, as well as the integration in autonomous vehicle models of the Carla simulator.

References

[1] M. Althoff, S. Maierhofer, and C. Pek, "Provably-correct and comfortable adaptive cruise control," *IEEE Transactions on Intelligent Vehicles*, 2020.

[2] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proceedings of the 1st Annual Conference on Robot Learning*, 2017.

[3] T. Korssen, V. Dolk, J. van de Mortel-Fronczak, M. Reniers, and M. Heemels, "Systematic model-based design and implementation of supervisors for advanced driver assistance systems," *IEEE Transactions on Intelligent Transportation Systems*, 2017.

[4] J. Krook, L. Svensson, Y. Li, L. Feng, and M. Fabian, "Design and formal verification of a safe stop supervisor for an automated vehicle," in *2019 International Conference on Robotics and Automation*, 2019.

[5] S. M. Loos, A. Platzer, and L. Nistor, "Adaptive cruise control: Hybrid, distributed, and now formally verified," in *International Symposium on Formal Methods*. Springer, 2011.

[6] P. Nilsson, O. Hussien, A. Balkan, Y. Chen, A. D. Ames, J. W. Grizzle, N. Ozay, H. Peng, and P. Tabuada, "Correct-by-construction adaptive cruise control: Two approaches," *IEEE Transactions on Control Systems Technology*, 2015.

[7] D. Nistér, H.-L. Lee, J. Ng, and Y. Wang, "The safety force field," *NVIDIA White Paper*, 2019.

[8] J. Park, D. Kim, Y. Yoon, H. Kim, and K. Yi, "Obstacle avoidance of autonomous vehicles based on model predictive control," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 2009.

[9] A. Rizaldi, F. Immler, B. Schürmann, and M. Althoff, "A formally verified motion planner for autonomous vehicles," in *Automated Technology for Verification and Analysis*. Springer, 2018.

[10] S. Sadraddini, S. Sivaranjani, V. Gupta, and C. Belta, "Provably safe cruise control of vehicular platoons," *IEEE Control Systems Letters*, 2017.

[11] S. Shalev-Shwartz, S. Shammah, and A. Shashua, "On a formal model of safe and scalable self-driving cars," 2017. [Online]. Available: http://arxiv.org/abs/1708.06374

[12] Q. Wang, D. Li, and J. Sifakis, "Safe and efficient collision avoidance control for autonomous vehicles," 2020. [Online]. Available: https://arxiv.org/abs/2008.04080