Will the Driver Seat Ever Be Empty?

Thierry Fraichard

Will the Driver Seat Ever Be Empty?

Thierry Fraichard

RESEARCH REPORT

N° 8493

March 2014

Project-Team PRIMA

ISSN 0249-6399 ISRN INRIA/RR--8493--FR+ENG

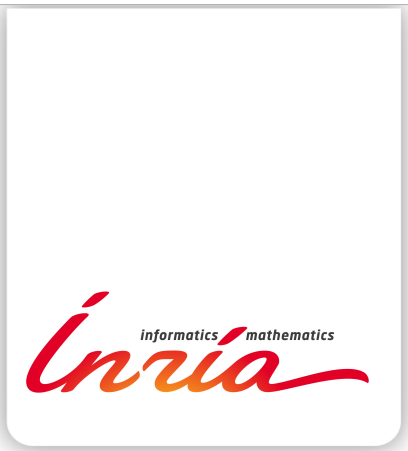

Will the Driver Seat Ever Be Empty?

Thierry Fraichard*

Project-Team PRIMA

Research Report n° 8493 — March 2014 — 14 pages

Abstract: Self-driving technologies have matured and improved to the point that, in the past few years, self-driving cars have been able to safely drive an impressive number of kilometers. It should be noted though that, in all cases, the driver seat was never empty: a human driver was behind the wheel, ready to take over whenever the situation dictated it. This is an interesting paradox since the point of a self-driving car is to remove the most unreliable part of the car, namely the human driver. So, the question naturally arises: will the driver seat ever be empty? Besides legal liability issues, the answer to that question may lie in our ability to improve the self-driving technologies to the point that the human driver can safely be removed from the driving loop altogether. However, things are not that simple. Motion safety, *i.e.* the ability to avoid collisions, is the critical aspect concerning self-driving cars and autonomous vehicles in general. Before letting self-driving cars transport people around (and move among them) in a truly autonomous way, it is crucial to assess their ability to avoid collision, and to seek to characterize the levels of motion safety that can be achieved and the conditions under which they can be guaranteed. All these issues are explored in this article.

Key-words: Mobile Robots; Motion Safety; Collision Avoidance;

* INRIA Grenoble-Rhône-Alpes, CNRS-LIG and University of Grenoble (FR).

RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Le siège conducteur sera t'il un jour vide?

Résumé : Les technologies de conduite automatique ont mûries et se sont améliorées au point que, au cours des dernières années, les voitures automatiques ont été en mesure de conduire en toute sécurité un nombre impressionnant de kilomètres. Il convient cependant de noter que, dans tous les cas, le siège du conducteur n'était jamais vide : un conducteur humain était au volant, prêt à prendre le relais dès que la situation dictée. C'est un paradoxe intéressant car le point d'une voiture automatique est d'enlever la partie la plus sensible de la voiture, à savoir le conducteur humain. Ainsi, la question se pose naturellement: le siège du conducteur sera t'il vide un jour? Outre les questions de responsabilité juridique, la réponse à cette question réside peut-être dans notre capacité à améliorer les technologies de la conduite automatique, au point que le pilote humain peut en toute sécurité être retiré de la boucle de conduite. Toutefois, les choses ne sont pas aussi simple que cela. La sécurité de mouvement, *i.e.* la capacité à éviter les collisions, est l'aspect critique à l'égard de voitures automatiques et les véhicules autonomes en général. Avant de laisser les voitures automatiques transporter des personnes (et se déplacer parmi eux) d'une manière réellement autonome, il est crucial d'évaluer leur capacité à éviter la collision, et de chercher à caractériser les niveaux de sécurité de mouvement qui peuvent être atteints et les conditions dans lesquelles elles peuvent être garanties. Toutes ces questions sont examinées dans cet article.

Mots-clés : Robots Mobiles; Sûreté de Mouvement; Evitement de Collision;

1 Introduction

One of the main stimuli behind the call for autonomous vehicles is safety. According to the World Health Organization, over 1.2 million people across the world die every year in road crashes, and between 20 and 50 millions are injured [1]. Because driver inattention and errors are responsible for most car crashes, it seems natural to strive to design self-driving cars. The first known attempt to build a self-driving car was in 1977 by the Tsukuba Mechanical Lab. in Japan. They demonstrated a car able to follow white markings and to reach speed up to 30 km/h on a dedicated circuit. Since then, a number of self-driving vehicles have been developed and tested, *e.g.* VITA I and II [2, 3], NavLab 5 [4], ARGO [5], VIAC [6] and Google's driverless cars [7]. These vehicles are reported to have safely driven an impressive number of kilometers[1] in different traffic conditions, *e.g.* highways, open roads, city streets. It should be noted though that, in all cases, the driver seat was never empty: a human driver was behind the wheel, ready to take over whenever the situation dictated it. This is an interesting paradox since the point of a self-driving car is to remove the most unreliable part of the car, namely the human driver. So, the question naturally arises: will the driver seat ever be empty? Will we ever witness truly self-driving cars on our roads? Legal liability issues may partly explain why the human driver has remained in the loop so far and why automotive manufacturers are primarily pushing to develop novel driving assistance systems, *e.g.* adaptive cruise control and pedestrian protection systems, instead of developing fully automated cars. Now, does the answer to the title question lies merely in our ability to improve the self-driving technologies to the point that the human driver can be removed from the driving loop altogether? Well, maybe, but things are not that simple....

The critical aspect concerning self-driving cars (and autonomous vehicles in general) is their *motion safety*, *i.e.* their ability to avoid collisions. Roboticists have long been aware of the motion safety issue and there is a rich literature on collision avoidance starting with the pioneering work of [8]. However, the accidents [9] that took place during the 2007 DARPA Urban Challenge (that called for self-driving cars to drive through an urban environment amidst human-driven vehicles) have showed that motion safety in the real world, *i.e.* an environment featuring moving obstacles whose future behaviour is uncertain, remains an open problem. Collisions happen for reasons that broadly fall into one of the following classes:

- Hardware failures, *e.g.* brake failure.
- Software bugs, *e.g.* truncation error.
- Perceptual errors, *i.e.* all the errors that are related to the sensor data processing systems of the vehicle and that result in the vehicle having an incorrect understanding of its environment (*e.g.* false negative).
- Reasoning errors, *i.e.* at a certain point a wrong decision is made.

In this article, we restrict ourselves to reasoning errors and look at motion safety solely from the decision-making point of view (which does not mean that the other aspects are not important, they are just as well). We feel that, before letting self-driving cars transport people around (and move among them) in a truly autonomous way, it is crucial to assess their ability to avoid collision, and to seek to characterize the levels of motion safety that can be achieved and the conditions under which they can be guaranteed. We focus on the case of dynamic environments, *i.e.* environments featuring moving obstacles (which is the case in most real world applications).

[1] In May 2012, Google's driverless cars were reported to have logged over 400 000 km without suffering any serious accidents.

Assuming that the autonomous vehicle at hand is working alright (from the hardware and software point of view) and has an accurate understanding of its current situation (no perceptual errors), we show that the presence of moving obstacles has a major impact when it comes to motion safety and that modeling choices determine what can (or cannot) be guaranteed and that ill-considered decision-making strategies are likely to yield collisions.

Following up on [10], we begin by exploring what motion safety is about. To that end, we use a toy scenario in order to gain insight into motion safety (§2). Then we turn to the *Inevitable Collision State* concept developed in [11] to further this analysis in a more formal framework (§3). Afterwards, we are ready to lay down a set of general motion safety rules whose violation is likely to yield collisions (§4). Next, because we have established the need to reason about the future evolution of the environment, we review the different classes of models of the future that are commonly used (§5). Finally, we discuss what levels of motion safety can (or cannot) be guaranteed and suggest ways to address the motion safety issue so that the driver will eventually move to the passenger seat (§6).

2 Case Study

In order to explore motion safety, we use a scenario dubbed *the compactor scenario* featuring two obstacles only (one fixed and one moving). As simple as this scenario may be, it provides insight into collision avoidance and helps in understanding key aspects related to motion safety.

2.1 Compactor Scenario

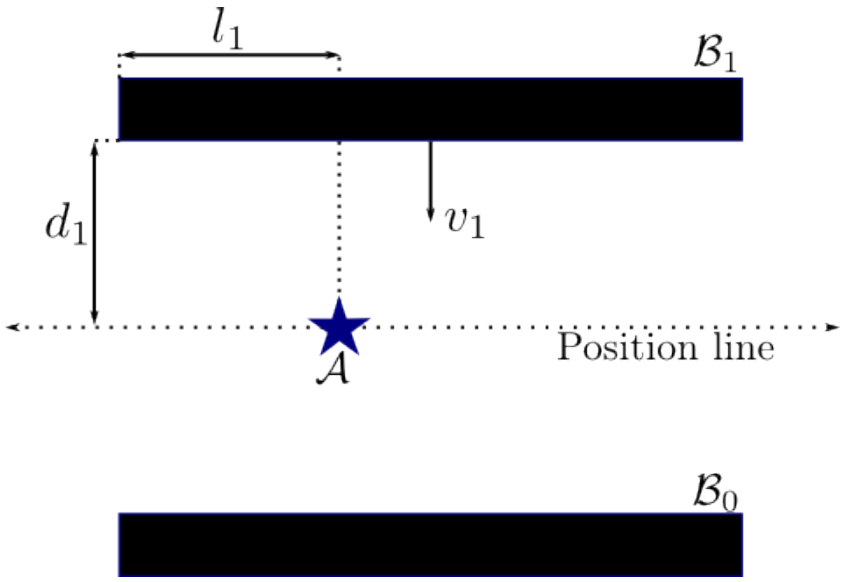

Figure 1: Compactor scenario: the plate $\mathcal{B}_1$ moves towards $\mathcal{B}_0$ at constant velocity v_1 until they meet.

Imagine a trash compactor or a car crusher, it can be modeled in 2D by two rectangular plates, one of them moving towards the other at constant velocity v_1 until they meet (Fig. 1). Let us put a robot $\mathcal{A}$ in the middle of the compactor. To avoid being crushed, $\mathcal{A}$ has to move

to the right or to the left until it exits the compactor. To further simplify the problem, $\mathcal{A}$ is treated like a 1D robot that moves along a horizontal line (henceforth called the *position line*). Assuming that $\mathcal{A}$ is a point robot which is velocity-controlled, a state of $\mathcal{A}$ is characterized by p that denotes the scalar position of $\mathcal{A}$ on the position line. It is finally assumed that the velocity of $\mathcal{A}$ is upper-bounded: $|v| \leq v_{\max}$.

2.2 Reasoning about the Future

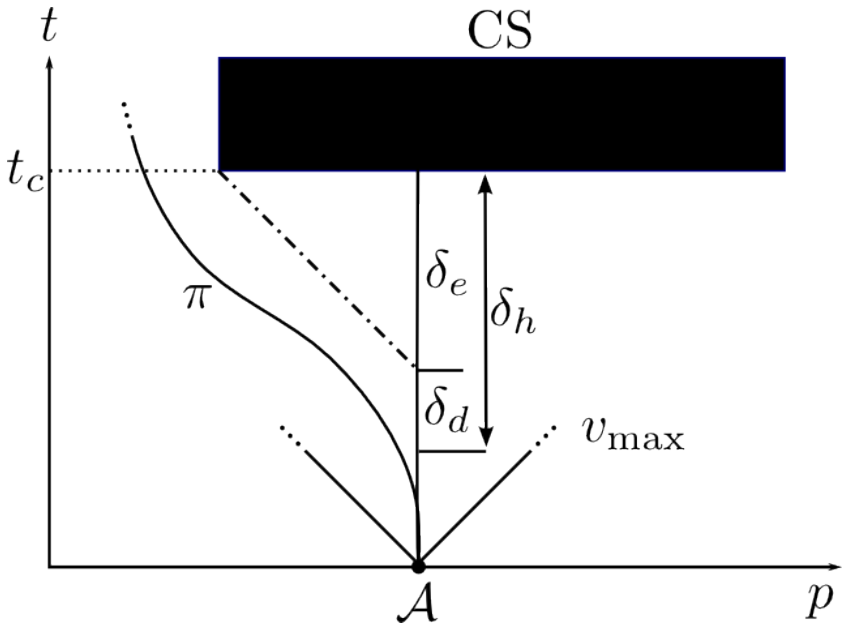

Figure 2: State×time space for the compactor scenario: the region CS is the set of state-times (p, t) where $\mathcal{A}$ is in collision with $\mathcal{B}_1$. The infinite cone whose apex is the current position of $\mathcal{A}$ and whose aperture is a function of $v_{\max}$ is the set of states that $\mathcal{A}$ can reach.

As trivial as it sounds, $\mathcal{A}$ needs to take into account the future motion of $\mathcal{B}_1$ in order to be aware of the upcoming collision risk. In the compactor scenario, the collision between $\mathcal{A}$ and $\mathcal{B}_1$ takes place at time $t_c = d_1/v_1$ where d_1 is the distance between $\mathcal{A}$ and $\mathcal{B}_1$. Since [12], it is generally acknowledged that *space-time* is the appropriate way to deal with moving obstacles. Adding the time dimension either to the configuration space [12] or the state space [13] of a robot allows to model the future evolution of the moving obstacles and therefore to reason about it, *e.g.* to plan a collision-free motion. In the compactor scenario, the state×time space of $\mathcal{A}$ consists of two-dimensional position and time. During its motion, $\mathcal{B}_1$ sweeps across the position line from time t_c onward. It yields a rectangular set of collision state-times (p, t) (the black rectangle labeled CS in Fig. 2). CS is a forbidden region that $\mathcal{A}$ must avoid. If needed be, the space-time model shows that if the future evolution of $\mathcal{B}_1$ is not taken into account, *e.g.* if $\mathcal{B}_1$ is treated like a fixed obstacle, the region CS does not appear in the space-time and $\mathcal{A}$ cannot be aware of the upcoming collision risk hence the importance of *modeling and reasoning about the future* evolution of the moving obstacles.

2.3 Limited Decision Time

To avoid being crushed, $\mathcal{A}$ has to move to the right or to the left until it exits the compactor. Let l_1 denote the distance to the nearest exit (on the left side in this case), the minimum time for $\mathcal{A}$ to escape the compactor is $\delta_e = l_1/v_{\max}$. $\mathcal{A}$ should therefore start moving to the left at least before time $t_c - \delta_e$ otherwise it does not have the time to escape. In other words, there is an upper bound on the time that $\mathcal{A}$ has in order to decide its future motion. Let δ_d denote the *decision time* of $\mathcal{A}$, δ_d must be strictly less than $t_c - \delta_e$.

2.4 Appropriate Time Horizon

Given the necessity to model the future evolution of the environment and reason about it, a question arises: with what *time horizon*, *i.e.* how far into the future should the modeling/reasoning go? In the compactor scenario, the answer is straightforward: the time horizon δ_h must be greater than $\delta_d + \delta_e$. Indeed, by setting δ_h to $\delta_d + \delta_e$, *i.e.* by considering the model of the future up until time $\delta_d + \delta_e$ only, $\mathcal{A}$ will become aware of the collision risk at time $t_c - (\delta_d + \delta_e)$ (see Fig. 2). It will leave $\mathcal{A}$ enough time to (a) decide that it should move to the left and (b) execute this motion. If $\delta_h < \delta_d + \delta_e$ then $\mathcal{A}$ is doomed.

As informal and intuitive as it may have appeared, the study of the compactor scenario has nonetheless brought to light three aspects that are important when it comes to motion safety. These aspects concern the time available to take a motion decision and the necessity to appropriately reason about the future evolution of the environment. In the next section, we will turn to the *Inevitable Collision State* concept developed in [11] in order to further this analysis in a more formal framework.

3 Inevitable Collision States

A space-time model such as Fig. 2 allows to model the future evolution of the environment and to represent the no-collision constraints in the form of forbidden regions. In this context, motion planning boils down to computing a feasible and collision-free trajectory π that drives $\mathcal{A}$ towards its goal. However, there is much more to motion safety than mere instantaneous no-collision. Imagine a car travelling very fast toward and a few meters away from a wall. Although the car is not in collision at the present time, it will crash regardless of any efforts to stop or steer. The concept of *Inevitable Collision States* (ICS) developed in [11] can be called upon to account for such a situation. An ICS is a state for which, no matter what the future trajectory of the robot is, a collision eventually occurs. Formally, an ICS is defined as follows:

Def. 1 (Inevitable Collision State) *a state s is an ICS if and only if $\forall \pi, \exists t \in [0, \infty] \mid \pi(s,t)$ is in collision.*

where π denotes a possible future trajectory of $\mathcal{A}$ and where $\pi(s,t)$ is the state reached by $\mathcal{A}$ at time t when starting from s. Similar to collision states (CS), ICS defines forbidden regions in the state-time space that must be avoided (it is actually a superset of CS). In the compactor scenario, the set of ICS is straightforward to characterize. It is the grey triangular region underneath the CS region (Fig. 3). Because of the upper-bound on $\mathcal{A}$'s velocity, as soon as $\mathcal{A}$ enters this grey region, it no longer has the time to exit from the compactor and it eventually collides with $\mathcal{B}_1$.

Note how the different aspects underlined in the previous section are present in the ICS concept. To begin with, Def. 1 shows that reasoning about the future is explicitly taken into

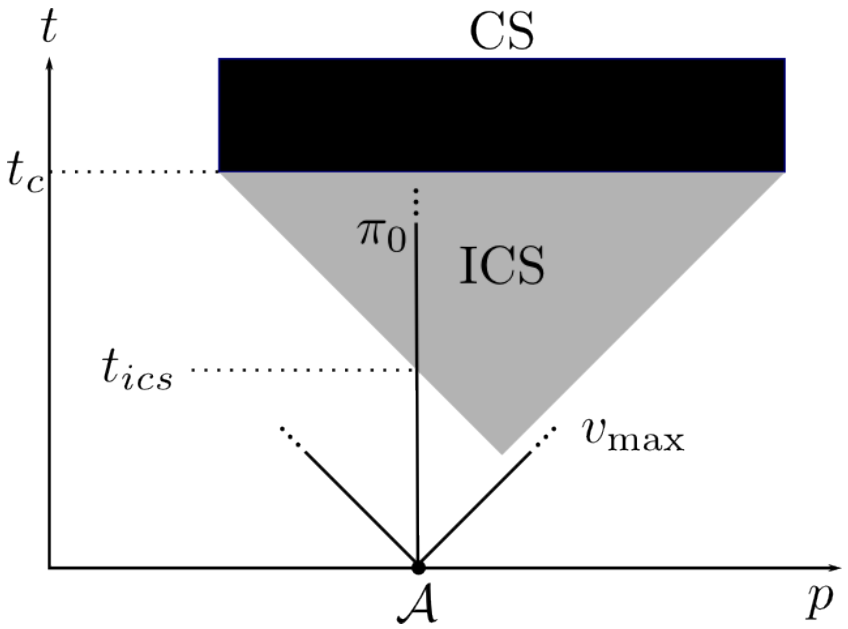

Figure 3: State×time space for the compactor scenario: it features both the set of collision states (CS) and inevitable collision states (ICS). If $\mathcal{A}$ stands still (*i.e.* following the trajectory π_0), it will enter the ICS region at time t_{ics}.

account (collision checking is done against the future position of the obstacles). Concerning the time horizon issue, it is implicitly taken care of due to the fact that the trajectories used to characterize the set of ICS have an infinite duration. In other words, ICS are defined for an infinite time horizon; it is therefore appropriate. Note however that ICS provides a theoretical way to compute the appropriate time horizon. Finally ICS also provides a straightforward way to compute the upper-bound on the decision time δ_d: Let π_0 denote the trajectory that $\mathcal{A}$ is currently executing. If π_0 drives $\mathcal{A}$ into the set of ICS at time t_{ics} then t_{ics} is the upper-bound on the decision time δ_d (Fig. 3).

The analysis of the ICS carried out in [11] has brought to light a property which is also important from the motion safety point of view: the set of ICS generated by a set of obstacles is not the union of the set of ICS generated by each obstacle independently. In other words, seeking to avoid collisions by considering each obstacle one by one may be a bad idea and lead the robot into situations where a collision becomes inevitable. It is important to consider the obstacles globally and not individually. To illustrate this, let us consider a variant of the compactor scenario featuring an additional moving plate $\mathcal{B}_2$ placed behind $\mathcal{B}_1$ with a slight offset to the left. $\mathcal{B}_2$ follows $\mathcal{B}_1$ with a delay. The resulting state×time space is depicted in Fig. 4. It features two set of collision states (one for each obstacle) and the two corresponding sets of ICS: ICS($\mathcal{B}_1$) and ICS($\mathcal{B}_2$). However, as soon as ICS($\mathcal{B}_1$) and ICS($\mathcal{B}_2$) overlap, the overall set of ICS is more than the mere union of ICS($\mathcal{B}_1$) and ICS($\mathcal{B}_2$). The dark grey region in Fig. 4 actually belongs to the set of ICS that must be avoided. If $\mathcal{A}$ enters this region, it will be able to avoid $\mathcal{B}_1$ alright but then, it will not be able to avoid $\mathcal{B}_2$ in spite of the fact that this dark grey region does not actually belong to ICS($\mathcal{B}_2$). From a motion safety point of view, it is therefore important to consider the obstacles globally.

The next section will summarize what we have learned so far about motion safety.

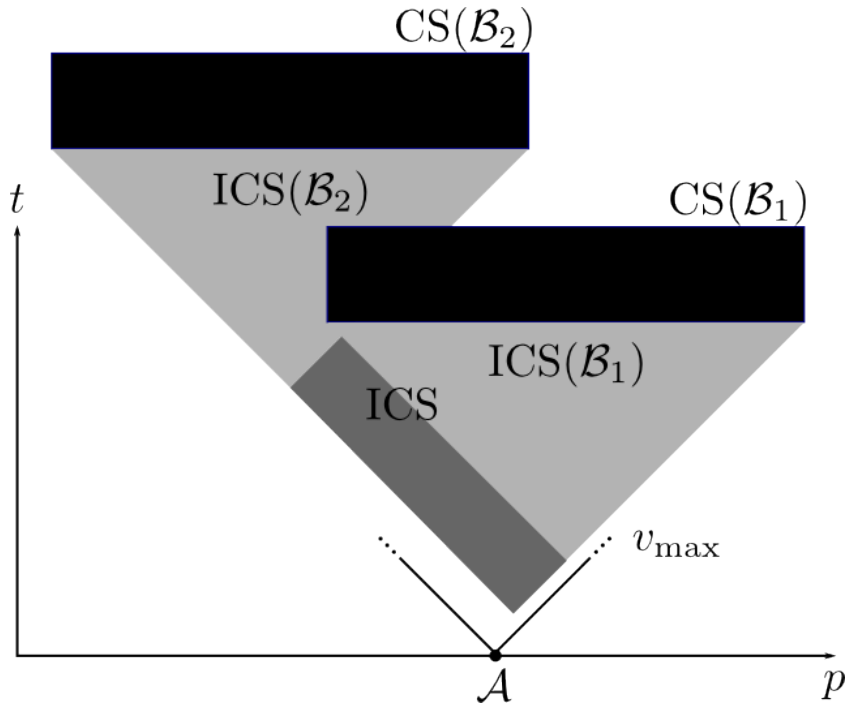

Figure 4: State×time space for the compactor scenario with two moving plates $\mathcal{B}_1$ and $\mathcal{B}_2$. Note how the actual set of ICS is more than the mere union of $\mathrm{ICS}(\mathcal{B}_1)$ and $\mathrm{ICS}(\mathcal{B}_1)$.

4 Motion Safety Rules

The insights on collision avoidance resulting from the study of the compactor scenario and the ICS concept are fairly intuitive and straightforward to express in two sentences:

> In a dynamic environment, one has a *limited time* only to make a motion decision. One has to *globally reason about the future* evolution of the environment and do so with an *appropriate time horizon*.

In other words, motion safety comprises four rules:

1. Decision time is upper-bounded.
2. Reasoning about the future is required.
3. Time horizon is lower-bounded.
4. Globally considering the obstacles is required.

Since rules 2 and 4 both concern the obstacles and their future evolution, they could arguably be merged together (would we then end up with the three laws of motion safety?). These rules may appear very abstract and general but, the important point is that if any one of these rules is violated then collisions are likely to happen (unless proven otherwise given the particulars of the situation at hand). Note how the first three rules above are all related to time. In a dynamic environment, *time is the critical aspect*. In the compactor scenario for instance, we have seen that the bounds on the decision time δ_d and the time horizon δ_h are:

$$\delta_d < t_c - \delta_e \text{ and } \delta_h \geq \delta_d + \delta_e \tag{1}$$

It is important to note that the bounds on δ_d and δ_h are largely determined by the current situation (through the position, size and velocity of $\mathcal{B}_1$). The bad news is that δ_d (resp. δ_h) can be arbitrarily large (resp. small). For instance, if $\mathcal{B}_1$ is close to $\mathcal{A}$ or moves fast then $\delta_d \longrightarrow 0$. Likewise, if $\mathcal{B}_1$ is very wide and very slow, *i.e.* $v_1 \longrightarrow 0$ and $l_1 \longrightarrow \infty$, then $\delta_h \longrightarrow \infty$. In other words, a robot can find itself in situations where (a) it has a very short time to decide its future motion, and/or (b) it must consider events that will happen very far into the future. These two contradicting constraints are very challenging from the decision-making point of view.

The ICS concept was initially investigated with the aim of designing navigation strategies for which collision avoidance could be formally guaranteed. The goal was to guarantee *absolute motion safety*, *i.e.* no collision will ever takes place whatever happens in the environment (hence the infinite time horizon). In a given situation, assuming that a model of the future is available up to the appropriate time horizon, the key to guaranteed motion safety is to characterize the corresponding ICS regions and stay away from them. To that end, one can use one of the numerous motion planning techniques currently available (see [14] for a recent survey of this topic). Provided that the decision time constraint can be satisfied, one ends up with a navigation strategy with proven collision-avoidance guarantee. The results reported in [15] and [16] demonstrate that is is possible *assuming that a complete model of the future is available*. As encouraging as these results are, we will see in the next section that things are not so rosy as soon as we are dealing with the real world.

5 Modeling the Future

The analysis above have stressed the necessity to model and reason about the future evolution of the environment. Building a space-time model such as that of Fig. 2 is in itself a challenge

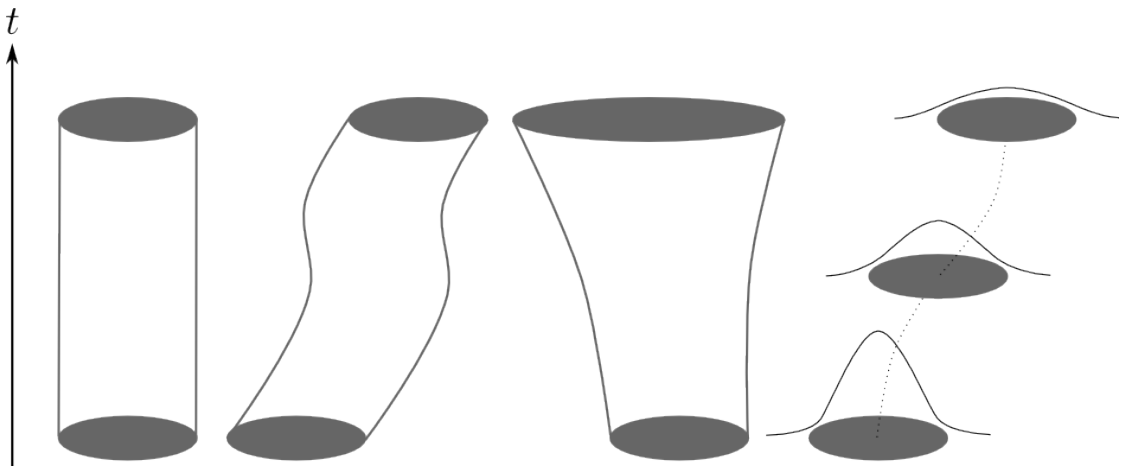

Figure 5: How to model the future? From left to right: deterministic (fixed, moving), conservative and probabilistic models for a disk obstacle.

inasmuch as, in most real-world situations, complete information about the environment and its future evolution is not available beforehand. To address this issue, a number of solutions have been proposed over the years. They yield models of the future that broadly fall into three classes: *deterministic*, *conservative* and *probabilistic* (Fig. 5). Let us now present and discuss these models from the motion safety point of view.

5.1 Deterministic Models

In such models, each obstacle is assigned a *nominal future motion* (Fig. 5-left). In certain situations, these nominal future motions are available beforehand, *e.g.* space applications. In most cases unfortunately, they must be predicted. The earliest deterministic models would consider every obstacle as a fixed obstacle. Later, with the progress in the area of the detection of moving obstacles, models of the future based on the prediction of the moving obstacles' future behaviour from their current state appeared. The prediction usually relies upon extrapolation, regression, or forward integration techniques. In other cases, sophisticated long-term motion prediction techniques have been proposed: they would either exploit the structure of the environment at hand or learn how the obstacles move in a given environment.

Given a deterministic model of the future, it is possible to develop a safe navigation strategy but its motion safety is only guaranteed with respect to the model of the future at hand which means that any discrepancy between the predicted future and the actual future voids the motion safety guarantee. From a motion safety point of view, deterministic models are useful as long as their prediction of the future evolution of the environment is reliable. Unfortunately, this reliability can decrease dramatically in the long-term. To address this issue, conservative models of the future have been proposed.

5.2 Conservative Models

In such models, the central idea is to consider all possible future motions of the environment's obstacles. Accordingly, each obstacle is assigned its *reachable set*, *i.e.* the set of positions it can potentially occupy in the future, to represent its future motion (Fig. 5-middle). The use of

conservative models solve the problem of the discrepancy between the predicted future and the actual future. Accordingly, it becomes possible to develop navigation strategies whose motion safety is guaranteed no matter what happens in the future (assuming that the reachable sets are accurately computed).

In theory, conservative models seem satisfactory from the motion safety point of view since they allow guaranteed collision avoidance. In practice however, the monotonous growth of the region potentially occupied by an obstacle is such that, eventually, the whole workspace will be potentially occupied by the obstacle. As a direct consequence of that, every state for the robot becomes an ICS since all trajectories eventually drive the robot to a collision state. Accordingly, any navigation strategy with proven collision-avoidance guarantee would fail to find a solution. To address this issue, probabilistic models of the future have been proposed.

5.3 Probabilistic models

In such models, the evolution of a moving obstacle is captured within a *stochastic transition function* and the tools used to predict the future behaviour of the moving obstacles are diverse, *e.g.* Kalman Filters, Hidden Markov Models and Monte Carlo Simulation. In this framework, the position of an obstacle at any given time is represented by an *occupancy probability density function* (Fig. 5-right). Probabilistic models are suited to represent the uncertainty that prevails in the real-world, in particular the uncertainty concerning the future behaviour of the moving obstacles. However, introducing probabilities clearly entails a major paradigm shift. So far, everything was black and white so to speak: collision *vs* no collision. When entering the realm of probabilities, everything turns grey and collision probabilities are in order. To address motion safety with probabilistic models of the future, [17] and [18] have both proposed probabilistic extension of the ICS concept. However, as sound as the probabilistic framework is, it cannot provide *strict motion safety guarantee*, strict in the sense that they can be established formally. Minimizing the collision risk is the only thing that can be done.

There is an interesting thing to note about probabilistic models of the future: as time passes by, the occupancy probability density function for a given obstacle diffuses, *i.e.* it flattens, and eventually cancels out. It leaves us then with the following paradox: with a conservative model, an obstacle is eventually everywhere while, with a probabilistic model, it is eventually nowhere, *i.e.* it vanishes (this property is exploited in [19] to compute an appropriate time horizon). From a motion safety point of view, both alternatives are not satisfactory.

6 Motion Safety Levels

At this point, it appears that absolute motion safety (in the sense that no collision will ever take place whatever happens in the environment) is impossible to guarantee in the real world (by that, we mean an environment featuring moving obstacles whose future behaviour is uncertain). The only way to attain absolute motion safety is to consider a conservative model of the future but we have seen how it renders the ICS concept ineffective. Today, most autonomous vehicles relies upon probabilistic modeling and reasoning to drive themselves. Probabilities are ideal to handle uncertainty but they will never allow strict motion safety guarantees. With respect to our title question, we may wonder then if humans will ever be ready to blissfully place their lives in the hands of a self-driving car whose sole asset it to minimize the collision risk. In an effort to improve the situation and to provide strict motion safety guarantees, we would like to advocate an alternative approach that can be summarized by the following motto:

Better guarantee less than guarantee nothing.

The idea is to settle for levels of motion safety that are weaker than absolute motion safety but that can be guaranteed. One example of such a weaker level of motion safety has been explored in [20]. It guarantees that, if a collision must take place, the robot at hand will be at rest. This motion safety level has been dubbed *passive motion safety*. It relies upon the definition of a new version of the ICS called Braking ICS. They are defined as states such that, whatever the future braking trajectory followed by the robot, a collision occurs before it is at rest. Formally, a Braking ICS is defined as follows:

Def. 2 (Braking ICS) *a state s is a Braking ICS if and only if* $\forall \pi_b, \exists t \in [0, t_b] \mid \pi(s,t)$ *is in collision.*

where π_b denotes a possible future braking trajectory of the robot and where t_b is the duration of π_b.

The key difference between ICS and Braking ICS is that, because Braking ICS consider braking trajectories only, it is possible to reason over a finite time horizon (function of the braking capabilities of the robot at hand) and therefore to use a conservative model of the future. Passive motion safety is readily obtained by avoiding Braking ICS at all times. The Braking ICS concept has been used to design a navigation scheme for a mobile robot with a limited field-of-view placed in an unknown dynamic environment. It has been formally established that this navigation scheme is *provably passively safe* in the sense that it is guaranteed that the robot will always stay away from Braking ICS no matter what happens in the environment. As limited as it may appear, passive motion safety is interesting for two reasons: (a) it allows to provide at least one form of motion safety guarantee in challenging scenarios (limited field-of-view for the robot, complete lack of knowledge about the future behaviour of the moving obstacles), and (b) if every moving obstacle in the environment enforces it then no collision will take place at all. In general, it could be interesting to explore more sophisticated levels of motion safety depending on the particulars of the navigation problem at hand. For instance, [21] suggested *passive friendly motion safety* that guarantees that, if a collision must take place, the robot will be at rest and the colliding obstacle could have had the time to stop or avoid the collision (if it wanted to). Such a motion safety level assume that the moving obstacles have perceptive and cognitive abilities, and are not hostile (which happens to be true in many situations).

7 Conclusion

Our goal in this article was to investigate if and how current self-driving technologies could be improved to the point that the human driver could safely be removed from the driving loop altogether. In our opinion, true self-driving will be achieved when it is be possible to design autonomous vehicles whose motion safety can be formally guaranteed. In the course of our investigation, we have brought to light the challenges imposed by the real world, *i.e.* an environment featuring moving obstacles whose future behaviour is uncertain. Challenges such that they rule out the possibility of ever guaranteeing *absolute motion safety* (in the sense that no collision will ever take place whatever happens in the environment). To make up for this harsh truth, we have advocated weaker safety levels as a possible answer to our initial question.

References

[1] "Global Status on Report on Road Safety: Time for Action," World Health Organization, Geneva (SW), Status Report, 2009.

[2] B. Ulmer, "VITA: an Autonomous Road Vehicle (ARV) for Collision Avoidance in Traffic," in *IEEE Intelligent Vehicles Symp.*, 1992.

[3] ——, "VITA II: Active Collision Avoidance in Real Traffic," in *IEEE Intelligent Vehicles Symp.*, 1994.

[4] T. Jochem, D. Pomerleau, B. Kumar, and J. Armstrong, "PANS: a Portable Navigation Platform," in *IEEE Intelligent Vehicles Symp.*, 1995.

[5] A. Broggi, M. Bertozzi, A. Fascioli, C. Guarino Lo Bianco, and A. Piazzi, "The ARGO Autonomous Vehicle's Vision and Control Systems," *Int. Journal of Intelligent Control and Systems*, vol. 3, no. 4, 1999.

[6] A. Broggi, P. Medici, E. Cardarelli, P. Cerri, A. Giacomazzo, and N. Finardi, "Development of the control system for the VISLAB Intercontinental Autonomous Challenge," in *IEEE Int. Conf. Intelligent Transportation Systems*, 2010.

[7] S. Thrun, "What we're driving at," The Official Google Blog, Oct. 2010. [Online]. Available: http://googleblog.blogspot.com/2010/10/what-were-driving-at.html

[8] H. Moravec, "Rover Visual Obstacle Avoidance," in *Int. Joint Conf. on Artificial Intelligence (IJCAI)*, Vancouver (CA), Aug. 1981.

[9] L. Fletcher, S. Teller, E. Olson, D. Moore, Y. Kuwata, J. How, J. Leonard, I. Miller, M. Campbell, D. Huttenlocher, A. Nathan, and F.-R. Kline, "The MIT–Cornell Collision and Why it Happened," *Journal of Field Robotics*, vol. 25, no. 10, Oct. 2008.

[10] T. Fraichard, "A Short Paper about Motion Safety," in *IEEE Int. Conf. on Robotics and Automation*, 2007.

[11] T. Fraichard and H. Asama, "Inevitable Collision States. A Step Towards Safer Robots?" *Advanced Robotics*, vol. 18, no. 10, 2004.

[12] M. Erdmann and T. Lozano-Perez, "On Multiple Moving Objects," *Algorithmica*, vol. 2, 1987.

[13] T. Fraichard, "Dynamic Trajectory Planning with Dynamic Constraints: a 'State-Time Space' Approach," in *IEEE-RSJ Int. Conf. on Intelligent Robots and Systems*, 1993.

[14] S. Lavalle, "Tutorial on Motion Planning for Dynamic Environments," IEEE Int. Conf. on Robotics and Automation, May 2012. [Online]. Available: http://msl.cs.uiuc.edu/~lavalle/icra12

[15] L. Martinez-Gomez and T. Fraichard, "An Efficient and Generic 2D Inevitable Collision State-Checker," in *IEEE-RSJ Int. Conf. on Intelligent Robots and Systems*, 2008.

[16] ——, "Collision Avoidance in Dynamic Environments: an ICS-Based Solution and Its Comparative Evaluation," in *IEEE Int. Conf. on Robotics and Automation*, 2009.

[17] A. Bautin, L. Martinez-Gomez, and T. Fraichard, "Inevitable Collision States: a Probabilistic Perspective," in *IEEE Int. Conf. on Robotics and Automation*, 2010.

[18] D. Althoff, M. Althoff, D. Wollherr, and M. Buss, "Probabilistic Collision State Checker for Crowded Environments," in *IEEE Int. Conf. on Robotics and Automation*, 2010.

[19] A. Kushleyev and M. Likhachev, "Time-Bounded Lattice for Efficient Planning in Dynamic Environments," in *IEEE Int. Conf. on Robotics and Automation*, 2009.

[20] S. Bouraine, T. Fraichard, and H. Salhi, "Provably Safe Navigation for Mobile Robots with Limited Field-of-Views in Dynamic Environments," *Autonomous Robots*, vol. 32, no. 3, 2012.

[21] K. Macek, D. Vasquez, T. Fraichard, and R. Siegwart, "Towards Safe Vehicle Navigation in Dynamic Urban Scenarios," *Automatika*, vol. 50, no. 3-4, Nov. 2009.

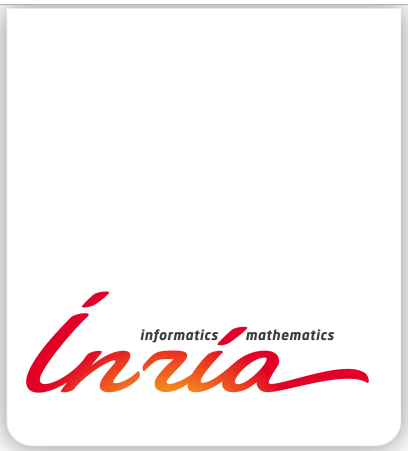

RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399