Achievable safety of driverless ground vehicles

Rodrigo Benenson
INRIA Rocquencourt / Mines Paris
78153 Le Chesnay Cedex, France
rodrigo.benenson@inria.fr

Thierry Fraichard
INRIA Rhônes-Alpes
38334 St Ismier Cedex, France
thierry.fraichard@inria.fr

Michel Parent
INRIA Rocquencourt
78153 Le Chesnay Cedex, France
michel.parent@inria.fr

***Abstract*—Safety is an important issue of driverless car. Yet, most current approaches fail to ensure safety even in a fully informed situation. In this paper we discuss how the safety criteria apply when the robot uses its on board sensors to evolve in a environment populated with static and moving obstacles. The sensors can only provide a partial and uncertain knowledge of the surroundings. We show that the usual safety notion does not apply for this relevant case and discuss which safety guarantees can be given and how to achieve them.**

I. Introduction

Driverless cars are an interesting robotics application. From a robotics perspective the main objective is to provide the robot car the capabilities to move alone from a point A to a point B in the city. Ensuring safety while moving the vehicle in a city is the core issue to be solved. Since most robot systems evolve in uninhabited areas or in highly controlled environments this issue appears as a novel problem to be addressed.

The problem of piloting in the city can be split in two: route planning and trajectory planning. Route planning worries about defining a route between points an A and B. Trajectory planning defines the precise motion allowing to traverse the sequence of streets defined by the route planner. Considering the wide variety of automotive navigation systems available in the market the problem of route planning (and re-planning) for a single vehicle can be considered solved[1]. This paper is concerned with the safety of trajectory planning in an environment such as populated cities.

Previous works discussed the need for a wide and dependable sensory coverage of the surroundings of the vehicle in order to provide enough information to avoid collisions [20], [19]. In [17], [22] the required detection ranges for typical roads manoeuvres are discussed (u-turns, overtakes, intersection crossing). Finally works like [9], [11], [3] proposes different approaches integrating perception, planning and control for driverless vehicles in urban environment.

In this paper we argue that most of existing works does not provide adequate safety guarantees for a driverless ground vehicle evolving in environments such as a city. We will show that the usual notion of safety based on inevitable collision states does not apply when the vehicle plans based on its sensors input. We argue that in such relevant case it is only possible to guarantee that the robot will not harm by action, we can not guarantee to be harmless through inaction.

[1]Deciding the route for a large group of vehicles considering the current road network status and the probable futures still being an interesting open problem.

In section II we provide a definition of safety and the criteria required to evaluate if an approach is safe or not. Section III defines the notions of uncertain and incomplete world model. Section IV explains the effect of uncertainty on safe planning and how to manage it. In section V we analyze the safety of a vehicle planning with an incomplete view of the world. Section VII discusses the effects of extra sensorial data on the safety guarantees. Finally section VIII offers some conclusions.

II. Safety

Mobiles robots, and especially driverless ground vehicles, are capable of harming them self and they surrounding environment (pedestrians, animals, other vehicles, the streets infrastructure, etc...). Motion safety relates on how to mitigate and, if possible, avoid harmful contact with the environment.

Definition 1: **(Safe motion)** *A robot is said to be safe if it can be guaranteed that its motion will not harm himself or its surroundings.*

For ground vehicles such harm can be created in three different ways:

- Traversing unfit terrain at inadequate speeds can harm the vehicle itself (running into water, holes, curbs, uneven or too inclined ground, etc...),
- Colliding with objects over the ground surface (pedestrians, animals, other vehicles, walls, poles, trash can, etc...),
- Pushing the vehicle out of its range of dynamic stability (overturn).

In any case harm will raise from setting the vehicle in the wrong state (position, speed, etc...) at a given time. In order to avoid harm we need to define the state of the vehicle in time.

Definition 2: **(Vehicle motion planning problem)** *Given an initial state $x(t_0)$, a measure of distance d to a desired state g and a world model w, define a feasible sequence of states in time that will generate a safe motion towards the desired state.*

The vehicle motion planning problem is an optimization problem with constraints. In our application the goal g will change in time, following a sequence of inflexion points of the roads defined during route planning. The *world model* is a representation of the world built from input data available to the robot (sensors, a priori information, communications, etc...). Since the planning relates to defining future states of the vehicle, the world model needs to provide information about the future state of the world. A sequence of states in time is named a *trajectory.* A trajectory is said *feasible* if it exists a sequence of commands that allows the robot to reach such states in time.

The safety of the motion does not only depends on the defined trajectory, but also on how the world model is built and used, and how the vehicle executes in the real world this trajectory. The safe motion of the robot is a propriety of the perception-planning-control trio.

A. Safety criteria

In order to analyze the safety of robotic systems three criteria have been proposed [6]:

1) Considering the motion of the robot
2) Considering the motion of the surrounding environment
3) Considering an infinite time horizon

The first criterion indicates that the limitation of the robot motion needs to be considered in the motion planning (maximum acceleration/deceleration, maximum speed, adherence constraints, etc...).

The second criterion indicates that the presence of moving and static obstacles around the vehicle has to be taken into account. The future position of surrounding objects needs to be considered in order to predict the free space available.

The last criterion indicates that, since the relative motion of the robot cannot be changed arbitrarily (first two criteria) at any point in the time the vehicle could be in course to an inevitable collision. Without any particular assumption on the real world checking collisions over a finite time horizon can not guarantee that the robot is not in such inevitable collision state [7]. It should be noted however that in many scenarios it is possible to define a finite set of verifications (computations) that will guarantee harmless behavior over an infinite time horizon [1], [12], [2].

When evolving in a environment populated by static and moving obstacles all of the traditional approaches to robot motion (nearest diagram, dynamic windows, velocity obstacles) fail to respect such criteria, and thus, fail to ensure harmless motion [6]. This occurs even when supposing a perfect knowledge of the present and future of the world.

As an example we can analyze the system proposed in [8]. They suppose that the world model built with they sensors is complete and that it correctly classify static obstacles and pedestrian. With each new sensor measurement the world model is updated and the vehicle updates its planned trajectory. The trajectory is generated from a set of predefined kinodynamically feasible control commands sequences. For each element of this set the collision with the world model is checked and a distance between the final state and the desired future state is measured. Trajectories leading to collision are disregarded and the sequence of commands leading nearest to the desired state (the goal) is selected. While announced as an approach for “secure driving in dynamic environments”, this method fails to consider the future motion of the pedestrians and provide no guarantee that in the next iteration at least one trajectory free of collision exists. Since it can not guarantee that no collision will occur, it fails to respect definition 1 and should be considered as non safe.

III. Building a world model from sensors

In order to solve the robot motion problem (see definition 2) we need to be able to evaluate two values. Let be $x(t)$ the state of the vehicle. Then $d(g(t), x(t))$ is the distance between the current state and the desired goal $g(t)$. This distance measure is used by the motion planner to search trajectories reaching the goal. The second value $h(x(t), w(t)) \in [0, 1]$ evaluates the harm caused by setting the vehicle in the state $x(t)$ given a prediction of the world state at time t. The notion of harm and its mapping to the function $h(x(t), w(t))$ are highly application dependent, for instance passing over a car may be fatal for a car, but harmless for a tank. This second value allows the motion planner to search for safe trajectories.

Based on the definition 1 we expect to be able to guarantee for a planned trajectory $\pi = \{ (x(t_1), t_1), (x(t_2), t_2), \ldots, (x(\infty), \infty) \}$ that $h(x(t_i), w(t_i)) = 0 \ \forall (x(t_i), t_i) \in \pi$.

Definition 3: **(Traversable space in time)** *A point p in space is considered traversable at time t if for any vehicle state $x(t)$ that covers the point p the harm value $h(x(t), w(t))$ is zero.*

Thus, one of the main concerns of a perception module, is estimating and predicting the traversable space in time.

An usual notion used when concerned on safe trajectory planning is the concept of *inevitable collisions state* (ICS).

Definition 4: **(Inevitable collision state)** *A state is considered in inevitable collision if all trajectories resulting from the application of any the possible commands sequence to this state, lead to a collision.*

In our case, since we are also concerned on the good of the vehicle itself, we will should speak if “inevitable damage state”. We will loosely use both terms as exchangeable. Strictly, we are concerned on avoiding reaching states where $h(x(t), w(t)) \neq 0$ (which, depending on the application, may include more that just collision states). In order to ensure safety the trajectory planner should, at least, only produce trajectories are free of inevitable damage states.

A. Incompleteness

When using on board sensors the robot is fundamentally limited to access only an incomplete and uncertain representation of the world. This means that some surrounding areas will be left unobserved (due to sensing range or occlusions, see figure 1), that the relative position of obstacles will be not necessarily known with precision (distribution of probability of collision over the space) and that the uncertainty in the prediction of future positions of moving obstacles will grow monotonously (since we do not know they future state). Beyond a certain point in the future nothing can be told about the traversable space.

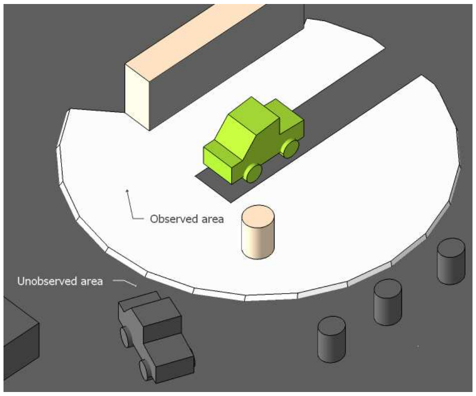

Figure 1. Illustration of observed and unobserved areas

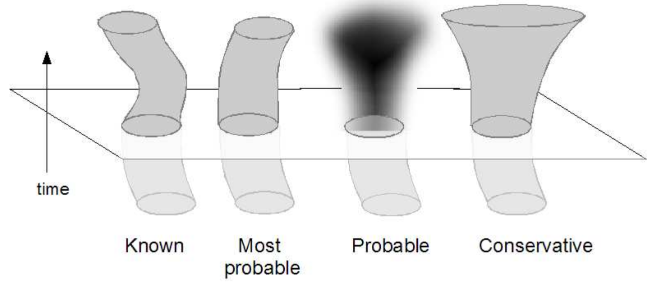

Figure 2. Predicting non traversable space in time. The region below the plane was observed, the region on top is predicted

Definition 5: **(Incomplete world model)** *A world model w is considered incomplete if $h(x(t), w(t))$ is not defined for every possible pair $(t, x(t))$.*
Since h is not defined everywhere, an incomplete world model can lead to incomplete or suboptimal plans.

B. Uncertainty

Sensors are corrupted by noise, thus relative distance measures to static obstacles will be uncertain. Also the state estimate of moving obstacles will be noisy. Even if a deterministic model existed for the moving obstacles, the initial uncertainty in the current state estimate will propagate in time. In many applications stochastic models are used to predict the motion of moving obstacles (e.g. motion of pedestrians or drivers) and thus position uncertainty grows in time. In an uncertain world model the traversability of a point p at an instant t becomes a probability value.

Definition 6: **(Uncertain world model)** *A world model w is considered uncertain when $h(x(t), w(t))$ is not available but the probability distribution $P(h(x(t), w(t)))$ is.*

IV. Uncertain world

Let us suppose by now that we have access to a complete but uncertain world model. As previously mentioned, this means that a given vehicle state in time is associated to a probability distribution of the harm value $P(h(x(t), w(t)))$. Since we seek for harmless motion (where $h(x(t), w(t)) = 0$) we are in particular interested in the value $P(h(x(t), w(t)) = 0)$.

If any moving obstacle exists in the world it future position uncertainty will grow monotonously in time. Supposing that the moving obstacles have a reachable space that cover ours (think of pedestrian and human drivers), some point in the future the whole world model will become uninformative: any traversable area may have become non traversable. What senses does it makes then to verify that no collision will occur up to infinity? Can we define an inevitable collision state in an uncertain world?

A. Cost function

An usual approach to safety consist on including it as part of a cost function over the possible trajectories. On an uncertain world model one could define such cost component using the integral of the vehicle trajectory over the probability distribution (in a spirit similar to [16]), as described in the equation 1.

$$safety\_cost = \sum_{(x(t_i), t_i) \in \pi} (P(h(x(t_i), w(t_i)) = 0) - 0.5) \tag{1}$$

Since defining a collision free trajectory up to infinity does not make sense, let us search for the trajectory to infinity with the lowest probability of collision. This approach is not suitable for multiple reasons. First we should remember that the primary aim of planning is reaching a specific point. The cost function will trade off collision risk with getting nearer to the goal, which undesirable if ensuring safety is a must. On the other hand, if collision risk has a very important height compared to reaching the goal, then generated trajectories will be unsatisfactory. When observing an empty area the safest trajectory is to stop as soon as possible. Using directly the cost function neither provides satisfactory safety nor satisfactory trajectories to reach the goal.

B. Probability threshold

In order to provide guarantees on safety, the probability distribution needs to be thresholded in order to define a binary function over space time (in a spirit similar to [23], for instance). Doing so is equivalent to providing a conservative estimation of the traversable space (see figure 2). We want to ensure that if some area in space time is predicted as traversable, it will be such in reality. By thresholding we ensure that the vehicle will always evolve in areas were a collision is considered highly unlikely and that the areas where a collision could occur will be left unvisited.

Definition 7: **(Conservative world model)** *A world model w' is considered a conservative approximation of w if and only if $h(x(t), w'(t)) = 0 \Rightarrow h(x(t), w(t)) = 0 \; \forall x(t) \forall t$.*
Thresholding $P(h(x(t), w(t)) = 0)$ allows to convert an uncertain world model into a certain world model by using

the function h' described in equation 2.

$$h'(h((x(t), w(t))) = h(x(t), w'(t)) = \begin{cases} 0, \ if \ P(h(x(t), w(t)) = 0) < p_{threshold}, \\ argmax \ P(h(x(t), w(t))), \ otherwise. \end{cases} \quad (2)$$

We expect h' (and thus, the selected threshold) to provide a conservative approximation of the real world (inaccessible through sensors).

When using a threshold over an incomplete and uncertain world model, it is likely that, over time, the predicted world collapses into a completely non traversable area. In this context the notion of inevitable collision states does not apply. This issue will be discussed in section V.

V. Incomplete world

Since the robot knows only a fraction of the world it is probable that it has not enough information to define a single definitive plan from its current state to the goal. Even if it did, as new information is acquired a better plan could be generated. This leads to the notion of *partial motion planning* [15], [14]. The denomination "partial" indicates that, unlike the usual approaches, the plan does not reach the goal (but, hopefully, leads towards it). Using a best effort approach the robot initially computes a partial plan to reach the goal. As the robots moves a new plan is computed to extend or enhance the previous one.

The criteria mentioned in section II imply that while the plan may be "partial" in the space dimension, it needs to be collision free over an infinite time horizon. To do so it is only necessary to ensure that every state of the plan is collision free and that the last state of the partial plan is not an inevitable collision state [13]. Verifying if a planned state will generate or not an inevitable collision is not a trivial problem, since we need to do verification over an infinite time horizon.

A. Static world

When the robot evolves in a static environment using a set of stop trajectories will provide safety, since we know that the vehicle will be able to stop before colliding and then no collision can occur (see figures 3 and 4). For this case, computing over a finite time horizon provides a guarantee over an infinite time horizon. Considering that a wall could appear at the frontier of the observed space, safety in an incomplete world is guaranteed if the unobserved space is considered as non traversable [1].

Swerving: Instead of stopping [21] proposed swerving the expected static obstacles, thus allowing higher speeds with the same sensors range. This approach makes strong assumptions on the maximum size and density of static obstacles, thus it is not suitable for highway or city like environments since the presence of a traffic jam (total road blockage) on a single direction road would generate a collision.

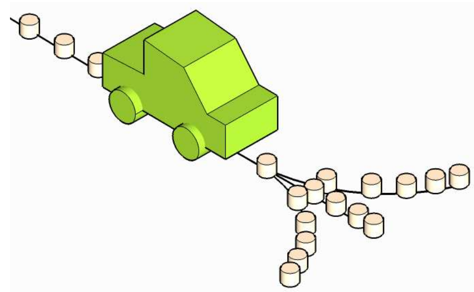

Figure 3. Trajectories resulting from some stopping commands for a vehicle. Dots are equidistant in time

B. Dynamic world

When the environment may include moving objects using stopping commands can not, by itself, ensure avoiding collisions.

Imitating manoeuvres: Ideally one would desire to have access to the present and future position of all the surrounding obstacles. When such information is available the notion of "imitating manoeuvre" has been proposed as a collision free trajectory over an infinite time [12]. However this approach is brittle since imitating a single moving obstacle could lead to colliding with another obstacle, and thus is not desirable in cluttered environments with arbitrary obstacles motions.

Finite time horizon: Knowing the motion from now to infinity of all the moving obstacles does not provide by itself a tractable way of detecting inevitable collisions states. Even if the vehicle can stop without having a collision with any moving or static obstacle, nothing prevents that a few seconds later a moving obstacle collide with the stopped vehicle. Defining a time horizon beyond which the vehicle is guaranteed to be collision free is an open problem in the general case. If it is possible to define a time t_l in the future beyond which no moving obstacle will approach the stopped vehicle, then verifying a collision of the stopping trajectory until the vehicle stops and beyond t_l is enough to ensure safety. This is approach used by [2] where it is supposed that the world is only populated by static obstacles and driverless vehicles, and the partial plan of every vehicle is perfectly known. At each re-planning step the trajectories of every vehicle are verified to be collision free with the previously planned ones and the re-planned ones. Since both the previous and the new partial plans finish with the vehicles stopped, it is possible to guarantee that the system is safe.

Best effort: Unless proof of the contrary an moving obstacles may be present at the frontier of the observed space. If we only have a partial knowledge of the surrounding environment a conservative world model needs to suppose the presence of moving obstacles at the frontiers of the unobserved area. Without information about the non observed moving obstacle we can not compute an avoidance trajectory. The best effort consists then in ensuring that the vehicle is able to stop without colliding. By itself, this does not ensure avoiding collisions, a moving obstacle could collide the vehicle immediately after it stops. Using stopping commands in this case can ensure that we will not harm by action, but only by inaction ("car on a railroad" scenario). This is the approach used in [13], [4].

In this approach we need to make assumptions about the moving objects that will appear, we need a model. This model predicts the possible (not the likely) presence of obstacles in time. The simplest possible model consist on "any obstacle, anywhere, in any direction" but with a bounded maximum speed. Depending on the specific application the model can become arbitrarily complex. For instance, in an urban scenario it could be reasonable to consider that no vehicle drives against the defined circulation directions. Doing so limits the possible appearance of moving obstacles and thus allows for planning higher speeds within the incomplete world model.

Theorem 1: If every moving obstacle uses a conservative world model and plans to stops before entering in a harmful state then no collision would occur and the overall system can be considered safe (see figure 4).

Proof: Let a and b be sensing moving objects. Both objects have a conservative world model with respect to the real world. The policy indicates that, since the world models are conservative, as soon as a detects b it has a trajectory allowing it to stop without colliding b or any other object. Conversely as soon as b detects a it already has a trajectory allowing to stop without colliding. Both a and b can stop without colliding, and the same applies to any couple of moving objects in the world. By induction since every moving object can stop without colliding, when one object is stopped, it is guaranteed that no other object will collide it, and thus it is a safe state over an infinite time horizon. ∎

If we ensure that the vehicle is able to stop given the possible appearance of an obstacle with a given maximum speed v_m and an obstacle appears coming faster than expected, then the vehicle will not be able to stop before the collision. Depending on the manoeuvrability of the vehicle, the clutter of the scene and the speed of the re-planning algorithm the vehicle may or may not be able to avoid the obstacles. *The safety guarantees are as good as the model* (safety depends not only on the planning algorithm but also on perception and control modules). If the predicted traversable space is not available, the guarantees that the motion of the vehicle will be harmless are lost.

VI. Safe planning

Based on the discussion of the previous sections we obtain the following definition.

Definition 8: **(Safe trajectory for driverless car)** *In an incomplete and uncertain world model of a dynamic environment a trajectory is considered safe if:*

- *Each state is collision free with respect to a conservative prediction of the traversable space in time,*
- *The sequence of states respects the vehicle capabilities,*
- *Its last state has speed zero.*

With the information available and without doing strong assumptions on the non observed areas, this is approach will ensure that no harm by action is done.

Please note that while the safe trajectory leads to stopping, it does not mean that the vehicle will necessarily stop. As the vehicle starts moving, the new sensors measurements will extend its world model. This allows it to replace the current trajectory (in execution) with a new safe trajectory. By doing so repeatedly the vehicle will drive continuously towards the goal and will not stop unless necessary (blocking obstacle). When using this approach the vehicle will automatically adapt its speed and behavior to the surrounding environment, considering both the observed and unobserved space. It will drive slowly in cluttered and uncertain situation, and faster on unlittered certain environments.

A. Unexpected events

Whatever are the passive or active measures taken, in the real world accidents will occur (e.g. "falling crane" scenario). Even using a conservative approach unexpected events will happen. As previously mentioned, if an area predicted as traversable is discovered as non traversable the harmless motion guarantee is lost. Maybe the vehicle will be able to avoid the collision. If not, it is important to be able to consider the cost of a collision. Crashing towards a trash can creates less harm than crashing a wall. Crashing an animal is better tolerated than crashing persons. If a collision becomes inevitable, it is desirable that the robot does not treat each possible smash equally.

This means that the perception system needs at least to localize the vehicle with respect to the planning goal (for planning) and to the plan currently in execution (for control purposes), to estimate the traversable space in time and to constantly monitor the cost of possible collisions.

When planning over a conservative world model, the planner expects that the robot will follow exactly the computed trajectory. Failing to do so would nullify any safety guarantee. This imply that the control module needs to provide a predictable bound on its tracking error. This bound is integrated in the planning stage to ensure that the trajectory is safe even with errors on the control. If the control unexpectedly fails to respect the predicted bound leave the robot in a situation equivalent to a violation of the conservative property of the world model, and a collision may or may not occur.

B. Safe perception-planning-control

Most of the previous works in the field will fail to respect the definition 8. For instance [23] proposed a planning method in dynamic environments able to deal with an incomplete and uncertain world model. However they do not take into account the unobserved area of the environment (non conservative estimation), and provide no guarantee that the vehicle will not collide if their probabilistic re-planning method fails.

In [5] a perception-planning duo was proposed to move in an incomplete and uncertain world model. The approach is similar to the one described here. However they algorithm restricts the movement of obstacle to constant speed vectors. Such approximation is non conservative over the obstacles considered in

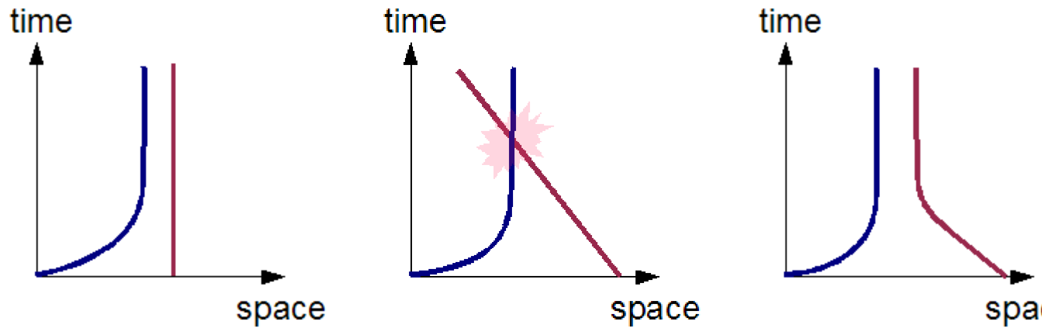

(a) stop in front of wall (b) stop in front of moving obstacle (c) both moving obstacles apply the same safety logic

Figure 4. Multiple scenarios of two obstacles in one space dimension. Blue and red lines are obstacles in time

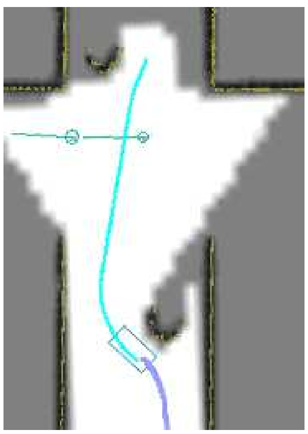

Figure 5. Example of safe planning in a perceived environment [4]

they work (cars intersection) and can be considered a violation of the safety criteria and the definition 8.

A solution consistent with definition 8 was presented in [4]. In that work, it is explained how the vehicle can efficiently estimate its movement, map static obstacles, detect and track moving obstacles, plan safe trajectories and execute them. All of it was integrated in a real world vehicle as a proof of concept prototype. Figure 5 presents an example of the output of the perception and planning of such algorithm.

More efficient and effective algorithms may be designed. Application specific models may provide more information about the future. Better sensors will provide larger observation areas and less uncertainty, allowing higher speeds. However we believe that safety can not be enhanced beyond the discussed limits. Real world safety is bounded by how good is the robot at estimating the possible futures.

VII. Complemented world

In the previous section we discussed safe motion of a robot using its on board sensors, a model of its sensor, a model of itself and a model of the existing moving objects. For the design of a mobile robot it is relevant to understand how other sources of information may affect the safety. When the robot completes the information observable from its point of view with other sources, we say it estimates a “complemented world model”.

A. Maps

There is the usual belief that high precision map will help robots navigation. These maps are of little use when concerned about safety. In populated environments such as the cities, if an area was seen as traversable a few hours ago, little can be said about if it is currently obstructed or not. Maps of the static environment are not able to provide reliable information of the future. Maps can be built to model the usual behavior of moving objects in a region [25]. This information can then be used to provide tighter predictions of the moving objects. Using this maps is still a delicate issue. Being built from observations it is hard to asses they completeness. Failing to predict a possible path for an obstacle, will lead to violation of the model and thus to a potential collision. Even worse, in the city, the behavior of pedestrians (for instance) not only depends on space but also on time (e.g. the Sunday street market) making it less likely to have a complete and reliable map of moving objects paths in space time.

B. Fixed path

A possible future for driverless vehicles in urban environment is the deployment of “immaterial tramways” [18], [24], [10]. Restricting the movement of the vehicle to a fixed path seems to be perceived by humans as a safer option.

In this configuration the planning is mainly concerned with speed control, a reduced set of commands ease the planning computation. However the needs for perception (building the world model) still being the same (with respect to range coverage or prediction capabilities). Knowing the path may provide a thin enhancement in the localization computation cost and in the collision cost estimate.

With a fixed path the exact same logic discussed on this paper applies. The safety issue is not relevantly modified.

C. Traffic rules

The use of traffic rules between humans seems to be a factor enhancing the safety on the roads. Supposing that other vehicles respect the rules, constraint they possible future movements and thus it is useful information to predict the traversable space. The use of traffic rules per se does not enhance the safety (since it depends on other vehicles respect of such rules), however it allows the vehicle to be less cautious in some situations (e.g. not slowing down at an intersection, because we have the priority).

D. Vehicle to vehicle communications

The use of vehicle to vehicle communication could enhance the safety in multiple ways. First, more measurements of the environment provides more information and thus probably larger observed areas and less uncertainty. Second, when the surrounding moving obstacles are driverless vehicles too, it would be possible to obtain they current plans, thus providing much tighter bounds than a naive worst case model [2].

If should be noted that the use of vehicle to vehicle communications requires to solve a relative positioning problem between the communicating vehicles. How to solve this problem in an unmodified city still being an open question. Even if a reasonable solution is provided, the relative positions will have a certain degree of uncertainty that will then be propagated over the transmitted data, reducing the benefits of the exchange.

VIII. Conclusion

Safety is a critical issue for driverless vehicles, however it still being a fuzzy aspect in many proposals. All of the classic planning methods are arguably unsafe. We reviewed the safety issue and how the safety apply to an incomplete and uncertain world model. While the usual notion of inevitable collision state has proved useful for analysis, it does not apply when the robot observes a dynamic world. We have shown that it is possible to ensure that within the limits of the used world model the vehicle will not actively harm. The final behavior obtained by a robot following the suggested approach will be as safe as humanly possible.

References

[1] K. E. Bekris and L. E. Kavraki. Greedy but safe replanning under kinodynamic constraints. pages 704–710, Rome, Italy, April 2007. IEEE press.

[2] K. E. Bekris, K. I. Tsianos, and L. E. Kavraki. A distributed protocol for safe real-time planning of communicating vehicles with second-order dynamics. Athens, Greece, Oct. 15-17 2007.

[3] Rodrigo Benenson, Stéphane Petti, Thierry Fraichard, and Michel Parent. Toward urban driverless vehicles. *International Journal of Vehicle Autonomous Systems, Special Issue on Advances in Autonomous Vehicle Technologies for Urban Environment*, 1(6):4 – 23, 2008.

[4] Rodrigo Benenson, Stéphane Petti, Michel Parent, and Thierry Fraichard. Integrating perception and planning for autonomous navigation of urban vehicles. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Beijing, China, 2006.

[5] Fulgenzi C., Spalanzani A., and Laugier C. Dynamic obstacle avoidance in uncertain environment combining pvos and occupancy grid. In *Proceedings of IEEE International Conference on Robotics and Automation, ICRA*, Rome, Italy, April 2007.

[6] Thierry Fraichard. A short paper about motion safety. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2007.

[7] Thierry Fraichard and Hajime Asama. Inevitable collision states - a step towards safer robots? *Advanced Robotics*, 18(10):1001–1024, 2004.

[8] S. Kolski, K. Macek, L. Spinello, and R Siegwart. Secure autonomous driving in dynamic environments: From object detection to safe driving. In *Proceedings of the Workshop on Safe Navigation in Open and Dynamic Environments: Applications to Autonomous Vehicles at the IEEE International Conference on Robotics and Systems*, San Diego, USA, 2007.

[9] Sascha Kolski, Dave Ferguson, Mario Bellino, and Roland Seigwart. Autonomous driving in structured and unstructured environments. In *IEEE Intelligent Vehicles Symposium*, pages 558–563, June 2006.

[10] Bouraoui L., Petti S., Laouiti A., Fraichard T., and Parent M. Cybercar cooperation for safe intersections. In *Proc. of the IEEE Int. Conf. on Intelligent Transportation Systems*, 2006.

[11] Christian Laugier, Stephane Petti, Dizan Alejandro Vasquez Govea, Manuel Yguel, Thierry Fraichard, and Olivier Aycard. Steps towards safe navigation in open and dynamic environments. In C. Laugier and R. Chatila, editors, *Autonomous Navigation in Dynamic Environments: Models and Algorithms*, Springer Tracts in Advanced Robotics Series (STAR). Springer, 2006.

[12] Rishikesh Parthasarathi and Thierry Fraichard. An inevitable collision state-checker for a car-like vehicle. In *IEEE Int. Conf. on Robotics and Automation*, 2007.

[13] S. Petti and Th. Fraichard. Safe motion planning in dynamic environments. In *IEEE-RSJ Int. Conf. on Intelligent Robots and Systems*, Edmonton, AB (CA), August 2005.

[14] Stéphane Petti. *Safe navigation within dynamic environments: a partial motion planning approach*. PhD thesis, Ecole des Mines de Paris, 2007.

[15] Stephane Petti and Thierry Fraichard. Partial motion planning framework for reactive planning within dynamic environments. In *Proc. of the IFAC/AAAI Int. Conf. on Informatics in Control, Automation and Robotics*, Barcelona (SP), September 2005.

[16] Roland Philippsen, Björn Jensen, and Roland Siegwart. Toward online probabilistic path replanning in dynamic environments. 2006.

[17] Scott Robertson. Motion safety for an autonomous vehicle race in an urban environment. In *Proceedings of the 2006 Australasian Conference on Robotics & Automation*, December 2006.

[18] E. Royer, M. Lhuillier, M. Dhome, and J.M. Lavest. Monocular vision for mobile robot localization and autonomous navigation. *International Journal of Computer Vision*, 74(3):237–260, 2007.

[19] Chuck Thorpe, Justin David Carlson, David Duggins, Jay Gowdy, Robert MacLachlan, Christoph Mertz, Arne Suppe, and Chieh-Chih Wang. Safe robot driving in cluttered environments. In *Proceedings of the 11th International Symposium of Robotics Research*, October 2003.

[20] Chuck Thorpe, Olivier Clatz, David Duggins, Jay Gowdy, Robert MacLachlan, James Ryan Miller, Christoph Mertz, Mel Siegel, Chieh-Chih Wang, and Teruko Yata. Dependable perception for robots. In *Proceedings of International Advanced Robotics Programme IEEE*, Seoul, Korea, May 2001. Robotics and Automation Society.

[21] Chris Urmson. Driving beyond stopping distance constraints. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, pages 1189–1194, October 2006.

[22] Christopher Urmson. *Navigation Regimes for Off-Road Autonomy*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, May 2005.

[23] Jur van den Berg, David Ferguson, and James Kuffner. Anytime path planning and replanning in dynamic environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2366 – 2371, May 2006.

[24] Jur van den Berg and Mark Overmars. Kinodynamic motion planning on roadmaps in dynamic environments. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems - IROS*, pages 4253 – 4258, 2007.

[25] A. D. Vasquez, Th. Fraichard, O. Aycard, and C. Laugier. Intentional motion on-line learning and prediction. In *Proc. of the Int. Conf. on Field and Service Robotics*, Port Douglas (AU), July 2005.