

A Short Paper about Safety

Thierry Fraichard

Inria Rhône-Alpes & Gravir Lab., Grenoble (FR)

<http://emotion.inrialpes.fr/fraichard>

Abstract—September 12, 2006

I. INTRODUCTION

Background and Motivation

Increasingly, mobile robotic systems are leaving the somewhat artificial world of the research laboratories. They are now trying to operate in the “real world”. Examples of such robotics systems can be found indoors or outdoors doing their best to carry out autonomously tasks as diverse as sweeping floors (*eg* Probotics Cye-SR, Gecko Carebot, iRobot Roomba), mowing lawns (*eg* Friendly Robotics RoboMower, Husqvarna Automower), moving goods in warehouses, factories and port terminals (*eg* Seegrid SmartTruck, BT Industries Autopilot, Frog Container Carriers), tour-guiding people in museums or shows (*eg* Rhino, Minerva, Robox, Rackham), helping people with disabilities (*eg* GuideCane, MAid), driving people around (*eg* Frog Parkshuttle and CyberCab), and even taking part in races (*cf* the Darpa Grand Challenge).

Designing an autonomous robotic system requires to solve a number of challenging problems in domains as different as perception, localisation, environment modelling, reasoning and decision-making, control, *etc.* However, whatever the robotic system and whatever the kind of tasks it is expected to carry out, at some point, it has to move. Motion is therefore a fundamental issue in Robotics. Motion safety is even more fundamental. As soon as the size and dynamics of a robotic system makes it potentially harmful for itself or its environment, the system should strive to avoid collision with the objects of its environment.

Now, with robotic systems designed to operate in the real world, among human beings in many cases, motion safety becomes critical. Before letting robotic tour-guides or automated cars operate autonomously, it is vital to assert their operational safety, *ie* their ability to avoid collision with the objects of their environment. The focus of this paper is precisely on motion safety with a special emphasis on safety in dynamic environments (since the real world, in most cases, contains moving objects: human beings, animals, vehicles or other robotic systems).

Contribution and Paper Outline

Motion autonomy is a long standing issue in mobile robotics. Since Shakey’s pioneering attempts at navigating around autonomously in the late sixties [1], the number and variety of autonomous navigation schemes that have been proposed is huge.

In general, these navigation schemes aims at fulfilling two key purposes: reaching a goal while avoiding collision with

the objects of the environment. When it comes to collision avoidance, once again, many collision avoidance schemes have been proposed (*cf* §III). Their aim of course is to ensure the robotic systems’ safety. However a careful analysis of these collision avoidance schemes shows that, in most cases and especially in dynamic environments, **safety is not guaranteed** (in the sense that it is relatively easy to find situations where collisions will take place). To some extent, this is due to the fact that safety is a concept that is taken for granted. In other words, the meaning of safety is never formally stated and, above all, the operational conditions of such collision avoidance schemes are seldom (if never) spelled out.

This paper is an attempt to change this state of affair. To begin with, three **safety criteria** are proposed. These criteria helps in understanding a number of key aspects related to the safety issue (§II). Then a number of popular existing collision avoidance schemes are evaluated with respect to these safety criteria. The following question is asked: is collision avoidance guaranteed, especially when they are used in dynamic environments? It turns out that most (if not all) collision avoidance schemes are not safe when used in dynamic environments (§III). The concept of **Inevitable Collision States** introduced in [2] is then called upon as an answer to the safety issue. An Inevitable Collision States (ICS) for a robotic system is a state for which, no matter what the future trajectory of the system is, a collision eventually occurs. Finally it is shown how the ICS concept embodies the three above-mentioned criteria and how it offers a theoretical answer to the safety issue (§IV).

II. SAFETY CRITERIA

Deciding one’s future course of action is a process that implies a certain amount of reasoning about the future: you decide now what you will do next. For a robotic system, the decision-making process is usually based on a model of the robotic system itself (usually given a priori), and a model of its environment. The environment model can combine a priori information (*eg* maps), sensor measurements, or computation results (*eg* prediction of the future). Besides the particulars of the decision-making process itself, the models that are considered have a direct impact on the decision which is taken.

This section examines three common sense criteria which, if violated, may put the actual robotic system into danger and yield a collision at some point in the future. These safety criteria are respectively related to the model of the robotic system, the model of the environment and the decision-making process. In all cases, it is assumed that a complete model of the environment is available, complete to the point that it also comprises information about the future motion of the moving

objects. The rationale behind this is that if the safety criteria applies with perfect information, it is expected them to be also relevant in the incomplete information case.

The safety criteria are all illustrated using the example of a one-dimensional point robot \mathcal{A} with double integrator dynamics (subject to velocity and acceleration bounds), and which is moving along a linear path. \mathcal{A} is characterised by its position p along the path, its velocity $|v| \leq v_{\max}$, and its acceleration $|a| \leq a_{\max}$.




Fig. 1. One-dimensional point robot example.




Fig. 2. The Position \times Time space of the one-dimensional point robot \mathcal{A} when it brakes down and stops (blue curve on the left). The green curve on the right corresponds to the fixed object \mathcal{B} .

A. Robotic System's Dynamics

Consider Fig.1 where a point object \mathcal{B} is at position p_b along \mathcal{A} 's path. \mathcal{A} should never decide to occupy position p_b since it would be in collision with \mathcal{B} .

Because of its dynamics, it takes \mathcal{A} a minimum time v/a_{\max} to slow down and stop. The distance travelled is:

$$d(v) = v^2/2a_{\max}. \quad (1)$$

If \mathcal{A} disregards its own dynamic characteristics, it could decide to occupy a position within the $[p_b - d(v), p_b[$ range since such positions are collision-free. Should this happen, \mathcal{A} would be in trouble because it would eventually crash into \mathcal{B} (no matter what it does in the future). Taking into account the dynamics of \mathcal{A} , the range $[p_b - d(v), p_b]$ becomes forbidden.

The same conclusion can be drawn by looking at the position \times time space of \mathcal{A} as depicted in Fig.2. It is assumed that, at time 0, \mathcal{A} is at position 0 with velocity v . When \mathcal{A} brakes down and stops at maximum deceleration, it traces a parabolic arc in the position \times time space. Once it has stopped, it traces a vertical line (the blue curve in Fig.2). The vertical

green line corresponds to the fixed object \mathcal{B} . When the distance between \mathcal{A} and \mathcal{B} is less than $d(v)$, the blue and green curve intersect each other meaning that a collision will occur.

This example illustrates the fact that, whenever a robotic system disregards its own dynamic characteristics, it may decide on a future course of action for which safety is not guaranteed and collision may take place eventually, hence the first safety criterion:

Safety Criterion 1: to decide its future motion, a robotic system should consider its **own dynamics**.




Fig. 3. The position \times time space of the one-dimensional point robot \mathcal{A} when it brakes down, stops and shifts in reverse until it reaches velocity v_b (blue curve on the left). The green curve on the right corresponds to the object \mathcal{B} moving at constant velocity v_b .

B. Environment Objects' Future Behaviour

Back to Fig.1 but assuming now that \mathcal{B} is moving to the left with a constant velocity $v_b \leq v_{\max}$. In this situation, it takes more than staying out of the $[p_b - d(v), p_b]$ position range to guarantee the safety of \mathcal{A} . Indeed, unless \mathcal{A} shifts in reverse until it reaches a velocity at least equal to v_b , a collision with \mathcal{B} will occur.

A straightforward analysis carried out in the Position \times Time space of \mathcal{A} allows to determine the range of forbidden positions in this case. Fig.3 depicts the Position \times Time space of \mathcal{A} . It is assumed that, at time 0, \mathcal{A} is at position 0 with velocity v . The blue curve represents the trajectory followed by \mathcal{A} when it slows down, stops and shifts in reverse at maximum acceleration until it reaches velocity v_b (it is made up of two parabolic arcs and a line segment). The green curve represents the trajectory followed by \mathcal{B} moving at constant velocity v_b . When the distance between \mathcal{A} and \mathcal{B} is less than $d(v, v_b)$, the blue and green curve intersect each other meaning that a collision will occur. It can be established that:

$$d(v, v_b) = (v + v_b)^2/2a_{\max}. \quad (2)$$

Taking into account both the dynamics of \mathcal{A} and the dynamics of \mathcal{B} , the range $[p_b - d(v, v_b), p_b]$ becomes forbidden (Fig.1).

This example illustrates the fact that, whenever a robotic system disregards the fact that an object is actually moving, it may decide on a future course of action for which safety is

not guaranteed and collision may take place eventually, hence the second safety criterion:

Safety Criterion 2: To decide its future motion, a robotic system should consider the **environment objects' future behaviour**.

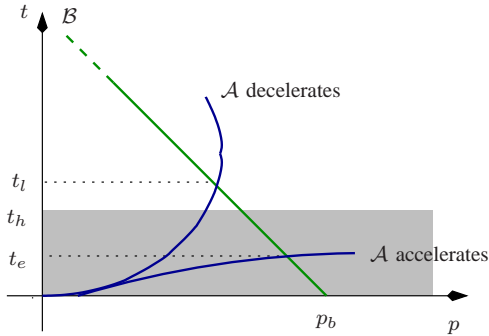


Fig. 4. Range of possible time to collision.

C. Time Horizon

If needed be, the previous two examples have revealed the fact that safety is not simply a matter of keeping the robotic system away from collision states. It is fundamental to keep the robotic system away from states that eventually yield a collision at some point in the future. To that end, reasoning about the future is required. Coming criterion explores the impact of the extent to which the future is explored.

Let us consider the situation where \mathcal{A} has to decide whether to occupy a position located within the range $[p_b - d(v, v_b), p_b]$ of positions forbidden *wrt* the moving object \mathcal{B} . In this case, collision does not take place right now, it happens later at a time instant that depends on the future behaviour of \mathcal{A} , *ie* whether it accelerates or decelerates (Fig.4). Let $[t_e, t_l]$ be the range of time instants when collision will take place. Both t_e and t_l are easily determined (roots of a quadratic equation). Assuming now that \mathcal{A} decides its future motion by restricting its reasoning to a finite time-horizon $t_h < t_l$ then \mathcal{A} may very well decide to occupy this position because, from its point of view, it is possible to avoid collision with \mathcal{B} (up to t_h , the decelerating trajectory is collision-free). Increasing the time horizon does not solve the problem because, in general, the future collision could happen at a time instant arbitrarily far away into the future (depending on the dynamic capabilities of \mathcal{A} and the future behaviour of \mathcal{B}).

In this respect, it is argued that, whenever a robotic system decides its future motion by restricting its reasoning to a finite time-horizon, collision may potentially happen beyond this time-horizon and safety cannot be guaranteed accordingly, hence the third safety criterion:

Safety Criterion 3: To decide its future motion, a robotic system should reason over an **infinite time-horizon**¹.

These three criteria are general and, depending on the circumstances, a navigation scheme may or may not have to

¹Or at least equal to the time required to reach the goal state, assuming the goal state is safe.

take them into account. For instance, when a robotic system is moving at slow speed, its dynamics can be ignored. Criterion 2 applies if the environment features moving objects. Likewise, in a static environment, a finite time-horizon corresponding to the robotic system's stopping time can safely be used. However, it should be emphasised that, in general (*ie* fast-moving robotic system, dynamic environment), all three criteria apply and violating either one of them may put the robotic system into danger and yield a collision at some point in the future.

III. ARE EXISTING ROBOTIC SYSTEMS SAFE?

As mentioned in the introduction, there has been a large number and variety of autonomous navigation schemes that have been proposed since the early days of mobile robotics. Such navigation schemes usually combines various perception, modelling, reasoning and control functions. A review of all these schemes is definitely out of the scope of this paper². This paper is interested on the collision avoidance components of such navigation schemes. Left aside are the “laboratory” robotic systems, the focus is on mobile robotics system that operate in real-world applications. Of particular interest are the robotics systems whose size and dynamics make them potentially harmful for themselves or their environment (given that, if a floor-sweeping robot such as Roomba bumps into a pet or a piece of furniture, it is really no big deal).



Fig. 5. Automated forklifts: Frog Palette Mover (left) and BT Industries Autopilot (right).




Fig. 6. Frog Container Carriers of the Europe Combined Terminal in the port of Rotterdam (NL).

The first family of mobile robotic systems considered are designed to transport goods (palettes, boxes, containers, *etc*) or people (for the most recent ones). They are commercial products and some of them have been operating for over a decade. They operate in warehouses, factories (Fig.5), port terminals (Fig.6), or road networks (Fig.7).

²The reader is referred to one of the books that address this topic, *eg* [3].



Fig. 7. Frog ParkShuttles in the business park of Rivium (NL).

They all operate along the same principle: a fleet of autonomous vehicles moves about a network of predetermined routes (Figs.6-left and 7-left). Their environments could be labelled as “protected” in the sense that unexpected objects are not supposed to be present on the route network (except for the odd pedestrian or fallen box). It explains why collision avoidance is very limited in such systems. Besides bumpers (that initiate a stopping manoeuvre upon contact), range finders are used to monitor the environment ahead of the vehicle. When an unexpected object is detected, a stopping manoeuvre is initiated (veering off the predetermined route is not an option). The stopping manoeuvre usually takes into account the dynamics of the vehicle but do not take into account the possible dynamicity of the object which means that collision can (and do) happen.

The second family of mobile robotics systems considered are human-size and designed to navigate among or interact with human beings. They are robotic tour-guides [4], [5], [6], [7], or automated wheelchairs [8]. They are not commercial products yet but have been deployed in real environments for a significant period of time. These robotic systems are interesting because their environments are extremely challenging with a possibly large number of moving objects interacting with one another and whose future behaviour is highly uncertain. The hardware and software architecture of these robotic systems are of course quite different. From the collision avoidance point of view however, they all rely upon one of the few popular collision avoidance approaches that have been proposed by the robotics community. These approaches are respectively the Nearness Diagram, the Dynamic Window and the Velocity Obstacle approaches. They are respectively reviewed in sections III-A, III-B and III-C. Section III-D analyses their performance in terms of safety.

A. Nearness Diagram

Rackham [6] is a tour-guide robot that operates in the Space-City museum³ in Toulouse, France (Fig.8-left). Collision avoidance is achieved thanks to the Nearness Diagram (ND) approach [9]. This reactive approach is similar in spirit to the earlier Vector Field Histogram approach [10]: a motion

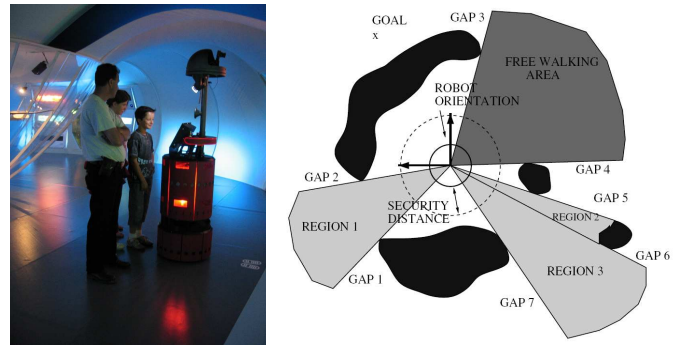


Fig. 8. The Rackham tour-guide robot (left) relies upon the Nearness Diagram approach for collision avoidance (right, source [9]).

direction is selected using a model of the environment surrounding the robotic system. This local model is built using sensor data and take the form of a polar distance histogram in which free angular sectors are computed (Fig.8-right). This approach has further been extended so as to allow the use of a global model, *ie* a map, of the environment [11]. The strength of ND primarily lies in the situation analysis which is carried out in order to select the motion direction. This situation analysis helps in reducing a number of problems that affects reactive navigation schemes, namely deadlocks and oscillations. Robust ND-based navigation have been demonstrated in very dense, cluttered and complex environments.

However, whatever its strength, it is important to note that the model that ND uses to take its motion decision is static: a moving object is considered as stationary. In other words, the second safety criterion (environment objects' future behaviour) is violated: safety in the presence of moving objects cannot be guaranteed.

B. Dynamic Window

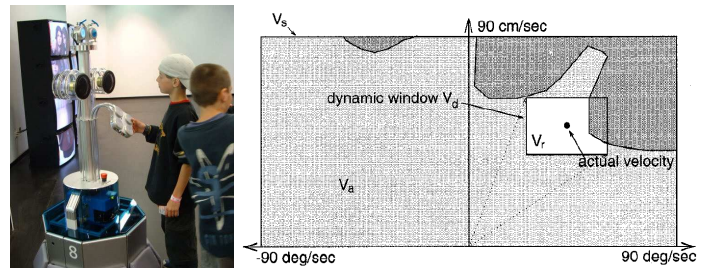



Fig. 9. The Minerva (top-left), Rhino (top-right) and Robox (bottom-left) tour-guide robots rely upon the Dynamic Window approach for collision avoidance (bottom right, source [5]).

³Http://www.cite-espace.com

Minerva [12], Rhino [4] and Robox [7] are robotic tour-guides that have operated for different time period in different places in the United States, Germany and Switzerland. Minerva was at the Smithsonian's National Museum of American History, Rhino at the German Museum in Bonn and Robox at the Swiss national exhibition Expo.02 (Fig.9). Collision avoidance is achieved thanks to the Dynamic Window (DW) approach [5]. This reactive approach operates in the velocity space of the robotic system considered. A velocity is admissible if it allows the vehicle to stop before hitting an object (Fig.9-right). An admissible velocity optimising a given criterion is selected at each time step. Robust DW-based navigation have been demonstrated at relatively high speeds (up to 1.0 m/s) in complex environments. This approach has further been extended so as to consider additional information about connectivity to the goal [13].

DW is superior to ND in the sense that the kinematics and dynamics properties of the robotic system considered are explicitly taken into account. The first safety criterion (robot's dynamics) is respected. However, like ND, the model that DW uses to take its motion decision is static: a moving object is considered as stationary. The second safety criterion is violated: safety in the presence of moving objects cannot be guaranteed.

C. Velocity Obstacle



Fig. 10. The MAid automated wheelchair (left) relies upon the Velocity Obstacle approach for collision avoidance (right, source [8]).

MAid [8] is an automated wheelchair that have been successfully tested in the concourse of the central station in Ulm (DE) and during the German exhibition Hannover Fair'98 (Fig.10-left). Collision avoidance is achieved thanks to the Velocity Obstacle (VO) approach [14]. This reactive approach also operates in the velocity space of the robotic system considered. Unlike DW, VO takes into account the velocity of the moving objects (assumed to be moving with a constant linear velocity). Each object yields a set of forbidden velocities whose shape is that of a cone (Fig.10-right). Should the robotic system select a forbidden velocity, it would collide with the moving object at a later time (possibly infinite) in the future. In practise, velocities yielding a collision occurring after a given time horizon are considered as admissible. Later, the approach was extended to consider arbitrary velocity profiles for the moving objects [15].

VO is superior to both DW and ND in the sense that both the kinematics/dynamics properties of the robotic system

considered and the moving objects' future behaviour are explicitly taken into account (safety criteria 1 and 2). However, the introduction of the time horizon violates the third safety criterion: safety in the presence of moving objects cannot be guaranteed.

D. Discussion

The different robotics systems presented above have all been up and running in crowded environments for a significant amount of time without any noticeable collision problems. Yet, it has been shown how they all violated one or several of the safety criteria introduced earlier! Does it mean that these criteria are meaningless? Not quite so. At the risk of being provocative, it is conjectured that the only reason why collision between these robotic systems and the people surrounding them did not happen is because and only because *people took care of the collision avoidance*. Had these robotic systems been placed among blind people for instance, collision could have happened. . .

to be completed. . .

REFERENCES

- [1] N. J. Nilsson, "Shakey the robot," AI Center, SRI International, Menlo Park, CA (US), Technical note 323, Apr. 1984.
- [2] T. Fraichard and H. Asama, "Inevitable collision states. a step towards safer robots?" *Advanced Robotics*, vol. 18, no. 10, pp. 1001–1024, 2004, [www].
- [3] I. R. Nourbaskhsh and R. Siegwart, *Introduction to Autonomous Mobile Robots*. MIT Press, 2004.
- [4] W. Burgard, A. Cremers, D. Fox, D. Hänel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun, "Experiences with an interactive museum tour-guide robot," *Artificial Intelligence*, vol. 114, no. 1-2, 2000.
- [5] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics and Automation Magazine*, vol. 4, no. 1, pp. 23–33, Mar. 1997.
- [6] A. Clodic, S. Fleury, R. Alami, M. Herrb, and R. Chatila, "Supervision and interaction: analysis of an autonomous tour-guide robot deployment," in *Proc. of the Int. Conf. on Advanced Robotics*, Seattle, WA (US), July 2005, pp. 725–732.
- [7] R. Philippsen and R. Siegwart, "Smooth and efficient obstacle avoidance for a tour-guide robot," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, Taiwan (TW), Sept. 2003.
- [8] E. Prassler, J. Scholz, and P. Fiorini, "A robotic wheelchair for crowded public environments," *IEEE Robotics and Automation Magazine*, vol. 8, no. 1, pp. 38–45, Mar. 2001.
- [9] J. Minguez and L. Montano, "Nearness diagram (ND) navigation: collision avoidance in troublesome scenarios," *IEEE Trans. on Robotics and Automation*, vol. 20, no. 1, pp. 45–59, Feb. 2004.
- [10] J. Borenstein and Y. Korem, "The vector field histogram — fast obstacle avoidance for mobile robots," *IEEE Trans. Robotics and Automation*, vol. 7, no. 3, pp. 278–288, June 1991.
- [11] J. Minguez, L. Montano, T. Siméon, and R. Alami, "Global nearness diagram navigation (GND)," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, Seoul (KR), May 2001, pp. 33–39.
- [12] S. Thrun, M. Bennewitz, W. Burgard, A. Cremers, F. Dellaert, D. Fox, D. Hänel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz, "Minerva: A second generation mobile tour-guide robot," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, Detroit, MI (US), May 1999.
- [13] O. Brock and O. Khatib, "High-speed navigation using the global dynamic window approach," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, Detroit, MI (US), May 1999, pp. 341–346.
- [14] P. Fiorini and Z. Shiller, "Motion planning in dynamic environments using velocity obstacles," *Int. Journal of Robotics Research*, vol. 17, no. 7, pp. 760–772, July 1998.
- [15] F. Large, C. Laugier, and Z. Shiller, "Navigation among moving obstacles using the NLVO : Principles and applications to intelligent vehicles," *Autonomous Robots Journal*, vol. 19, no. 2, pp. 159–171, September 2005.