

Safety Strategy for Autonomous Systems

Rami Debouk, Ph.D.; Barbara Czerny, Ph.D.; Joseph D' Ambrosio, Ph.D.;

Jeffrey Joyce, Ph.D.; Critical Systems Labs. Inc.; Vancouver, British Columbia, Canada

Keywords: Autonomous Driving Conditions, Driver Attentiveness, Safety Strategy

Abstract

In November 2007, an un-manned Chevrolet Tahoe – nicknamed “Boss” - successfully “drove” itself through a 60-mile urban course to win the prestigious Defense Advanced Research Projects Agency (DARPA) Urban Challenge competition. “This competition significantly advanced the understanding of what is needed to make driverless vehicles a reality, as GM continues to reinvent the automobile”, said Larry Burns, then GM’s Vice President for Research and Development.

Today’s vehicles already feature an emerging family of electronic driver-assist technologies aimed at reducing driver errors that can result in crashes. These technologies are the first steps in developing driverless or autonomous driven vehicles and they need to be carefully and systematically analyzed to meet and exceed all government and manufacturer requirements.

In performing the safety analysis of these emerging electronics-enabled autonomous driving vehicles, factors such as driver attentiveness and external conditions for autonomous driving (e.g., visibility of lane markings) need to be considered in addition to the vehicle’s potential malfunctioning behavior. This paper discusses some of these factors and proposes a safety strategy to account for these factors’ impact during the safety analysis.

Introduction

The development of dependable motion control systems has played a key role in the design of vehicles in the automotive domain. In early vehicles, motion control systems such as steering or braking systems were designed using mechanical components and operated independently of one another. Over the last two decades, vehicle design has evolved to include the use of electronics and software in these motion control systems. For example, steering systems are now often realized by electric power steering systems, comprised of mechanical components and electronics that include software. Currently, individual features of a motion control system still operate mostly independent from one another and mostly activate based on driver input. For example, a decision by a feature to request longitudinal acceleration may be made independently of a request for an adjustment to steering by another feature.

Recently, the automotive industry has begun introducing new features in vehicles that may autonomously command its motion control systems, such as crash mitigation braking. These features when providing commands to the vehicle motion control systems, may direct them to take actions independent of the driver under some conditions. This represents a departure from previous motion control systems, in that they may now operate independent of the driver and multiple motion control systems may activate simultaneously based on commands received from one or more features. Thus, the introduction of highly integrated features can be seen as an enabler for autonomous driving. Indeed, a few prototype or development autonomous vehicles have been introduced lately.

In 2007, The Chevrolet Tahoe “Boss” [1] won the DARPA Urban Challenge competition that required self-driven vehicles to navigate an urban course (including merging traffic, intersections, and stop signs) for 60 miles in less than 6 hours. The Chevrolet “Boss” was equipped with a suite of sensors to achieve front, rear, and side sensing. Data collected from these sensors was fused and fed into control algorithms intended for vehicle path planning, threat assessment, and collision avoidance maneuvers. The output of these algorithms was commands to the vehicle motion control systems to enable its self-navigation on the urban course.

More recently, other autonomous prototype vehicles were introduced by the VisLab at the University of Parma in Italy in the middle of 2010 [2]. The VisLab vehicles (four of them) completed 13,000 km from Italy to China with the lead vehicle driven autonomously the majority of the time. Similar to “Boss”, these vehicles were equipped with

a suite of sensors (cameras and scanners) and used on-board control algorithms to fuse the data, assess threats, and plan the path to be followed. Limited, human interventions were needed for the lead vehicle to define the route and intervene in critical situations.

Late 2010, Google deployed a fleet of autonomous vehicles in California [3]. These vehicles are equipped with radar sensors, laser range finders and map data to plan and control the motion of the vehicle. The automated vehicles are manned by trained operators and software developers that are capable of maintaining safe operation of the vehicle. Google has logged more than 14,000 miles on these vehicles.

While these research prototypes are impressive demonstrations of the future possibilities of a dramatic increase to the level of computer-based automation in a passenger vehicle, the prospect of selling a production vehicle with an autonomous driving capability for use on public roads depends on solutions to a number of technical challenges.

With the addition of complexity necessary for autonomous driving, many challenges arise with respect to the safety analysis of these vehicles. Component malfunctions and failures, as well as the interaction amongst the different functions / features that are achieving autonomous driving need to be considered as potential causes for hazards. These interactions may result or lead to unintended behaviors of the vehicles [4,5]. Therefore, additional focus needs to be given to the application of the system safety process in that respect. In addition to considering failures and functions / features interaction, additional challenges related to the enabling conditions for autonomous driving need to be addressed: these conditions should be defined and accounted for when performing the hazard analyses as their violation could be a cause for potential hazards. Finally, the fact that the driver may not be controlling all functions / features of the vehicle needs to be factored into the analysis.

In this paper we present a safety strategy to account for the challenging issues raised above. We propose the strategy to account for the impact of physical failures, violation of environmental conditions, and driver inattentiveness during an autonomous driving scenario. In following the proposed strategy, we address items that may be identified by the system safety analysis activities and demonstrate one way to address the challenges identified above.

This paper is organized as follows. We provide a comparison of the technical challenges of autonomous driving with autonomous control in other modes of transportation, e.g., the autopilot of a commercial passenger aircraft. Next, we discuss requirements for autonomous driving. The proposed safety strategy is then presented and demonstrated through an example.. Afterwards, we provide some design strategies for autonomous driving. Finally, we summarize our results.

Autonomous Operation in other Modes of Transportation

Various forms of autonomous control have been incorporated in other modes of transportation for many decades. The first autopilot in an aircraft was developed by Sperry Corporation at the dawn of manned flight nearly one hundred years ago. Since then, the level of automation available to pilots has increased dramatically and forever changed the man-machine relationship. Not only has technology advanced, but also the readiness of society to accept some forms of automation has increased. Looking beyond aviation, we can find other examples of autonomous control that are no longer remarkable. Autopilots for ships followed soon after aircraft. Driverless trains have existed for nearly a half century and are now used in many metro transportation systems of the world.

Given the history of technical innovation in the transportation sector, the introduction of autonomous road vehicles may seem long overdue. However, the technical challenges of providing road vehicles with autonomous driving capability involve a number of considerations that distinguish this capability from autonomous control in other modes of transportation such as flight, maritime and rail.

One such consideration is the separation in time between a normal situation and an undesirable situation. When flying in controlled airspace, an aircraft is required to be separated from other aircrafts, terrain and other dangers by a distance or amount of time that would normally provide the flight crew with ample opportunity to take control of the aircraft when the autopilot has failed or is behaving in a manner deemed unsafe by the flight crew. This contrasts sharply with the separation in time between a normal situation and an undesirable situation in the case of autonomous driving, which might be measured in seconds rather than minutes. For example, a failure to adjust the steering angle of a vehicle entering a sharp turn at highway speed could cause a lane departure in seconds.

A related consideration is the fact that other modes of transportation such as flight (in controlled airspace) and rail are systems (or systems of systems) that include rigorously enforced safety margins (such as margins for wing and body stress loads, and proximity restrictions in aircraft). At least in the case of flight and rail systems, the separation rules are actively monitored and infractions of these rules are generally regarded as serious matters. There are rules for the operation of passenger vehicles on public roads, such as speed limits. However, these rules are less conservative than the separation rules for other modes of transportation. In addition, the frequency at which rules for passenger vehicles are violated is higher than for other modes of transportation.

Another consideration is the extent to which equipment is maintained. In general, commercial airlines adhere to stringent maintenance procedures. In contrast, no assumptions can be made that individual vehicle owners will provide maintenance other than that required to keep the vehicle on the road. Poor maintenance is not a new source of risk with respect to the safety of passenger vehicles. However, we suggest that it is one of the considerations that distinguish the technical challenges of autonomous driving for passenger vehicles from autopilots in airplane and ships, and from driverless train operation.

A fourth consideration concerns the skills, training and general fitness (e.g., physical and mental) of the human operators. Pilots regularly spend time in flight simulators performing training scenarios that include failures of various kinds of automation including potential failures of the autopilot, for example. Thus, pilots are well prepared to react to unexpected situations such as a failure of the autopilot. Typically, passenger vehicle drivers do not receive specialized training and thus are not trained to deal with unexpected driving scenarios.

Another consideration is the extent to which the environment is predictable in forms of autonomous control in domains other than road vehicles. Commercial passenger aircraft operate in a highly predictable environment. Air traffic controllers provide pilots flying in controlled airspace with clearances that specify planned maneuvers (e.g. route, speed, altitude) that leave little room for uncertainty about the trajectory of the flight compared to what may be encountered by an autonomously driven vehicle travelling along a busy freeway, for example. A high degree of predictability also exists for driverless train operation. The signaling system typically moves trains according to a pre-determined schedule which, if everything “runs smoothly”, is assured to be conflict free. In contrast, the task of driving a passenger vehicle along a busy freeway typically involves frequent adjustments to speed in response to nearby traffic, and maneuvers such as changing lanes to avoid potential imminent problems, e.g., avoiding an obstacle in the current lane.

Thus, the technical challenges of providing road vehicles with autonomous driving capability involve a number of considerations that distinguish this capability from autonomous control in other modes of transportation. There may be additional considerations as well that add to the technical challenges faced by the automotive industry in the development of autonomous driving capability for passenger vehicles. However, the considerations discussed in this paper should clarify why the development of such a capability in road vehicles is not simply a minor adaptation of existing technology already available for other modes of transportation.

Requirements on Autonomous Driving

Since the introduction of the automobile in the late 19th century, the human driver has been more or less in charge of all of the driving tasks. Driving tasks include (but are not limited to) the following:

- starting the vehicle
- driving the vehicle including
 - braking
 - shifting
 - accelerating
 - steering

- maintaining situational awareness, that is knowledge about the surrounding environment
- communicating with other vehicles
 - signaling for a turn
 - honking the horn
- stopping the vehicle
- parking the vehicle

Electronics and software based control systems have been introduced into the automotive domain for some time now. These systems support the driver in providing assistance while s/he performs the driving tasks. For instance, an automatic starter turns the vehicle engine on, however the vehicle cannot be started until the driver inserts and turns the ignition key on or presses the start button. Another example is that of adaptive cruise control whereby the driver is supported by automatic acceleration and/or deceleration depending on the traffic situation to maintain a constant headway with respect to the lead vehicle. In both examples, it is understood that since these are assist features, the driver is responsible for maintaining control of the vehicle at all times and for remaining aware of the surrounding environment as well (we will use later the term “engaged in driving” to refer to the awareness of the surrounding environment). The following table provides a list of some of these electronic-based driver assist features and some of the driving tasks they support or augment.

Table 1— Driver Assist Features

Assist Feature	Cruise Control	Adaptive Cruise Control	Lane Keeping / Lane Centering	Park Assist	Collision Mitigation Braking
Driving Task					
Braking	X	X		X	X
Shifting	X	X		X	X
Accelerating	X	X		X	
Steering			X	X	

Many driver assist features, such as cruise control, can be categorized as “fail silent” features, i.e., in the event of a failure, they should produce no output (for example, no commands to the braking actuators). The possibility that a fail silent feature might not perform its intended function does not make a vehicle equipped with this feature any less safe than a vehicle not equipped with this feature. An example from Table 1 is Collision Mitigation Braking (CMB). CMB is engaged when the feature determines that the driver of the vehicle will be unable to stop in time to avoid a collision. CMB may not be able to avoid a potential crash in all cases, however, it will be able to reduce the energy that results from a collision. The fact that a vehicle is equipped with CMB does not decrease the responsibility of the driver to operate the vehicle in a safe manner – for example, maintaining adequate separation from a preceding vehicle in the same lane. In addition, the installation of CMB does not decrease the responsibility of the driver to be prepared to respond with braking and possibly steering in a situation where there is an imminent danger of a collision.

The categorization of driver assist features, such as those listed in Table 1, may need to be re-considered when they are integrated into an autonomous driving capability. Depending on certain assumptions about how the driver is expected to use this capability, it may not be acceptable to categorize some driver assist features as fail-silent. For example, it might be assumed that appropriate use of the autonomous driving capability does not require the driver to maintain a grip on the hand-wheel at all times, if this capability includes the Lane Keep Assist feature (described below). If this particular feature were to experience a fault, the driver may not take back control of the steering quickly enough to avoid any potential harm. Instead, it may be necessary to categorize the driver assist feature as a “fail-operational” feature – which means that, ideally, it will continue to fulfill its intended purpose at least until the

driver can be reasonably expected to take back control of the steering. The fail-operational behavior may be degraded with respect to normal, full performance operation.

However, not every feature that is integrated into an autonomous driving capability needs to be fail-operational. Some features may remain fail-silent even if they are an integral part of the autonomous driving capability.

In the rest of this paper, we will use the example of a lane keep assist (LKA) / lane departure warning (LDW) feature to illustrate our proposed strategy. We define the LKA feature as one that reads the visible lane markers of certain roads and applies slight torque to the steering wheel to help keep the car in the center of the lane. It is not a collision-avoidance system and is not a substitute for safe and attentive driving. To help keep the driver from drifting out of his or her lane, LDW is defined to give an audible warning when the outside perimeter of the vehicle approaches a lane boundary within some specified distance. After the warning, LKA takes corrective steering action to help keep the vehicle in the current lane.

Although LKA is not a substitute for safe and attentive driving, the fact that it has the capability to assist the driver with steering the vehicle to help maintain the lane (alongside the availability of the adaptive cruise control feature) represents a first step in introducing autonomous driving into the automotive domain. Not only do new driver models need to be defined in order to understand what the responsibilities of the human driver are with respect to the driving tasks with respect to autonomous driving, but also the conditions under which autonomous driving is allowed need to be specified. A driver model should include assumptions on the time required by the driver to recognize their surroundings and to intervene in the case of a required manual maneuver. If intervention of the driver is required, then a driving monitoring system should be considered to make sure the driver is engaged in driving even though the vehicle is being driven autonomously. The specification of the autonomous driving conditions would dictate when these features are allowed to operate.

Approach

The factors discussed above need to be carefully looked at when considering driving scenarios related to these emerging electronics-enabled autonomous driving vehicles and determining the safety strategy to apply. Driver attentiveness and conditions for autonomous driving need to be considered in addition to the vehicle's potential malfunctioning behavior when identifying and categorizing potential hazards. We propose a model for a safety controller of a vehicle with autonomous driving capabilities and describe a safety strategy for application to autonomous road vehicles. A model for a safety controller used in the proposed safety strategy is shown in Figure 1.

Figure 1 schematically illustrates a model of a controller managing subsystems in charge of providing autonomous driving capabilities. The controller monitors and analyzes the specific conditions necessary for autonomous driving and helps ensure fault handling and degradation when one or more of the specific conditions necessary for autonomous driving no longer exist, or will no longer exist at some time in the future. The controller is composed of the following elements:

- Supervisor
- Sub-systems 1 ... n
- Human Machine Interface (HMI)
- Driver Monitoring System (DMS)
- Autonomous Driving Monitor (ADM)

The supervisor has supervisory control over the multiple sub-systems, the HMI, the DMS, and the ADM. The status of these systems is communicated to the supervisor and it generates action commands to them based upon the monitored information that is communicated between them. The multiple subsystems locally monitor diagnostic conditions indicative of systematic failures/internal faults that can impede their performance, and where use of fault handling may be required. The monitored diagnostic conditions are continuously communicated to the supervisor

that determines the occurrence (or possibility of occurrence) of a failure. The supervisor is configured to directly fix a detected failure within one of the subsystems, if it is feasible to do so. The HMI contains numerous mechanisms, including but not limited to, buttons, voice activation, etc. for the operator to activate and cancel operation of the system when the operator so desires. This is important as the driver is assumed to be able to override the autonomous driving functions at any time. The DMS includes monitoring the attentiveness of the driver. Such a system helps ensure the driver is still engaged in driving - although the vehicle is being driven autonomously – just in case the driver is required to take control of the vehicle. The ADM primarily monitors the external conditions necessary for autonomous driving. For instance, some weather conditions or some road types may not be suitable for autonomous driving and such a monitor would not allow the autonomous driving features to engage if requested by the driver. The ADM communicates with the supervisor that analyzes the monitored information from and determines if one of the specific conditions for autonomous driving no longer exists or may expire in the near future. As mentioned above, the supervisor issues command actions based upon the monitored information. For instance, the supervisor can warn an inattentive driver to take control of the vehicle through the HMI, if required.

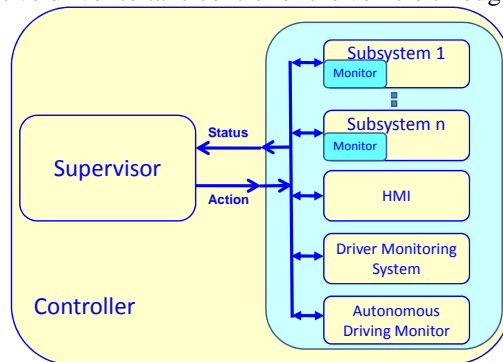


Figure 1 — Model for Autonomous Driving Vehicle

We propose a safety strategy for managing the various systems and subsystems that considers some of the factors that play a role in determining the operational state of a vehicle driven in an autonomous mode in the event of undesirable conditions/events. Figure 2 provides a graphical depiction of an example safety strategy for autonomous driving in the presence of the factors of interest discussed in this paper (e.g., physical failures, driver attentiveness, autonomous driving conditions). The modes of operation shown in Figure 2 represent example modes that may be included in autonomous driving systems. These modes include autonomous driving disabled and autonomous driving enabled. Autonomous driving enabled is a superstate that contains the superstate’s autonomous driving modes and monitoring activities. The autonomous driving modes superstate contains the following sub-states: full autonomous driving, warning mode autonomous driving, degraded autonomous driving, and autonomous driving shutdown. The monitoring activities superstate includes the driver monitoring, and autonomous driving monitor substates. The two superstates within the autonomous driving enabled superstate operate in parallel.

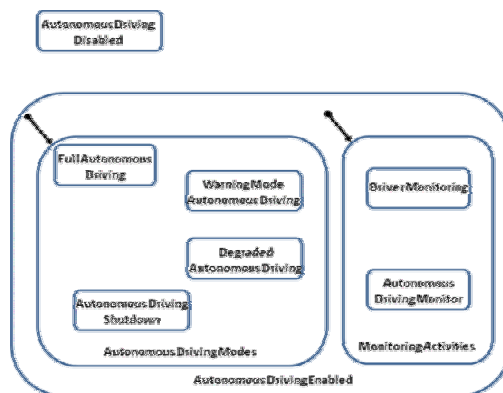


Figure 2 — Example Safety Strategy for Autonomous Driving

In our example, when the vehicle is in one of the autonomous driving modes, the monitoring activities are continuously monitoring the environment around the vehicle, potential fault conditions, and the driver attentiveness.

If the autonomous driving monitor determines that conditions necessary for autonomous driving are no longer satisfied, or will expire in the near future, the system will transition to a warning mode autonomous driving. In warning mode autonomous driving, the driver is notified that they should take control of the vehicle. At the same time, the autonomous driving system is examining the environment around the vehicle and planning potential evasive or fail-safe maneuvers to bring the vehicle to a desired state should the driver not take control of the vehicle after a specified amount of time. The warnings in warning mode autonomous driving consist of various levels depending on the response from the driver and depending on the urgency of the request. For example, if the driver does not respond to the initial warnings, a heightened warning will be issued to the driver. If the driver does not respond to the heightened warnings, the system will operate in a degraded mode to bring the vehicle to a desired state that was planned based on information the system gathered on the environment surrounding the vehicle when the warnings to the driver were being issued. Once the system has brought the vehicle to the desired state, autonomous driving is shutdown. If the driver responds and takes control of the vehicle, then autonomous driving is disabled.

There are various levels of concerns that may trigger a warning to the driver. These may be that the vehicle has detected that environmental conditions necessary for autonomous driving are about to expire (for example, lane markings will disappear), or a fault has occurred that precludes full operation of autonomous driving. In the event of a fault, there are various levels of faults that elicit various levels of response from the system. In some cases, sufficient time exists to wait through several warnings for a response from the driver. In other cases, very little time exists to wait for a response from the driver from a single warning, and in some cases, the warning state may be bypassed. In any of these cases, the vehicle may enter a degraded state sooner in order to begin to maneuver the vehicle to a desirable state so that autonomous driving can be shutdown as soon as is feasible.

It is possible that the reasons for the warnings being given are no longer present before autonomous driving is degraded, before autonomous driving is shutdown, or before the driver responds to the warnings and takes control of the vehicle. For example, if the system detected that conditions necessary for autonomous driving (e.g., lane markings) are not present, and subsequently begins to warn the driver, but then detects that the conditions are present again (e.g., lane markings reappear) such that autonomous driving can continue unaffected, then the driver warnings would cease, and the vehicle would return to full autonomous driving mode.

Figure 3 superimposes an example lane keeping assist feature onto the safety strategy model used for autonomous driving, to demonstrate the complexities added to a system when autonomous driving is available. For LKA without autonomous driving capabilities, the modes include Lane Keeping Disabled and the superstate Lane Keeping Enabled (LKE). Within the LKE superstate, are the two parallel superstates, Lane Keeping Modes (LKM) and Monitoring Activities (MA). In the absence of autonomous driving, there are only three substates within the LKM and MA superstates; Lane Keeping, Warning Mode, and Lane Marker Monitor.

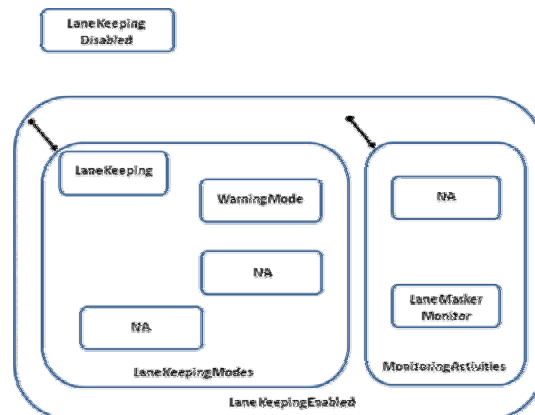


Figure 3 — Example Safety Strategy for LKA

Figure 4 shows the LKA feature as part of an autonomous driving feature. The figure indicates the amount of complexity that is added in moving from individual driver assist systems integrated into the vehicle, but working

independently, to an autonomous driving feature that contains a number of driver assist systems integrated and working collectively as part of an autonomous driving feature.

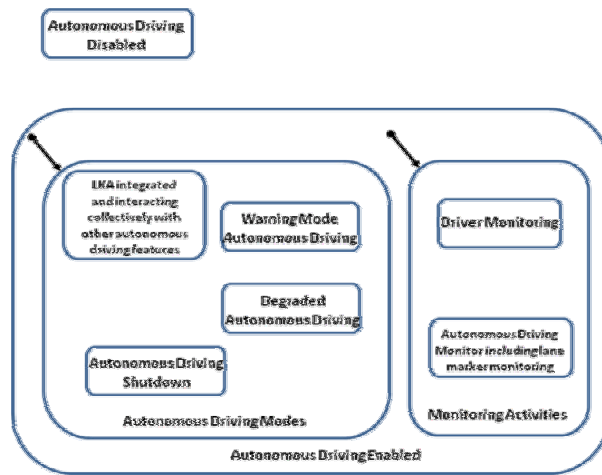


Figure 4 — Example Safety Strategy for LKA as part of Autonomous Driving

Design Strategies

In this section we provide an overview of some potential design strategies for implementing the high integrity hardware and software needed for autonomous driving vehicle systems. The approach described here is based upon using fail safe / fail silent control units as a building block to construct fail operational units implementing the needed autonomous driving functionality.

A fail safe / fail silent control unit is intended to provide high integrity control for systems in which it is acceptable to turn the control unit off when a fault occurs. Thus, the off state must be an acceptable safe state, and the manner in which the control unit transitions from operation to the off state when a fault occurs must itself be safe as well. Additional modes of operation, providing partial, reduced functionality, may be available depending on the specifics of the control application being implemented.

Figure 5 shows a typical hardware design pattern for implementing a fail safe control unit. The primary controller executes the autonomous driving control function. The “safety” controller is responsible for monitoring the “primary” controller, and thus provides an independent mechanism to shut the control unit down when a fault occurs. Both control units monitor the system for faults, and if a fault is detected, outputs are disabled. High integrity sensor inputs may be provided by using redundant sensing modules. For autonomous driving applications, this may take the form of diverse sensing technologies. For example, information from a forward-looking camera, one or more forward-looking radars, and one or more vehicle-to-vehicle or vehicle-to-infrastructure wireless communication sensors can be integrated to provide the vehicle with a high integrity view of objects in front of the vehicle.

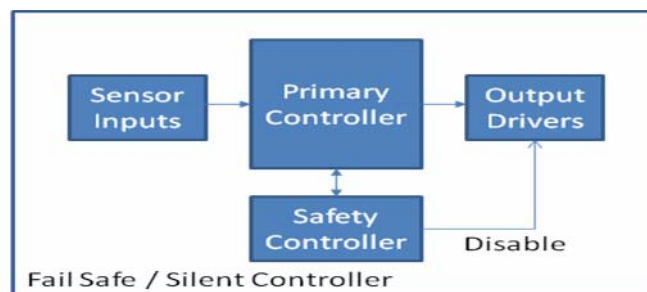


Figure 5 — Fail Safe / Fail Silent Control Unit




Figure 6 — Software Strategy for Fail Safe / Fail Silent Control Unit

Figure 6 above shows a typical software approach for detecting faults with the system. The “level 1” software represents the control application, which provides the primary functionality of the system. This software runs on the “primary” control unit. The “level 2” software implements diagnostic software focused on monitoring the application behavior (i.e., the “level 1” software), to detect unexpected errors. Typically the level 2 software runs on the primary control unit. The “Level 3” software implements diagnostics focused on detecting random hardware failures of both control units. The level 3 software is split between the two control units, and is used to perform diagnostic self checks as well as diagnostic cross checks between the two control units.

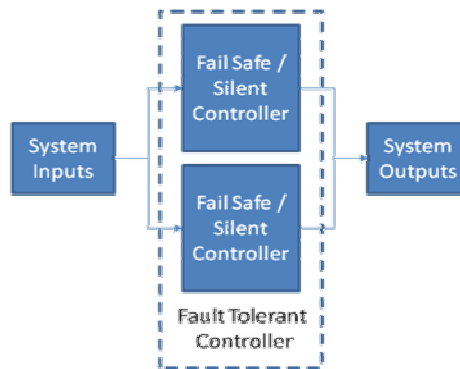


Figure 7 — Fail Operational Unit

Figure 7 shows a typical hardware design pattern for implementing a failure operational unit that can be used for autonomous driving applications. The fail operational unit is composed from two fail safe / silent units, in which the behavior that must fail operationally is implemented in both control units. When a fault occurs in one of the fail safe / silent units, it detects the fault and shuts down. However, the remaining fail silent unit continues to operate, providing the required system behavior. Note that system inputs are provided to both fail safe / silent control units, and that the system outputs must be designed so as to be able to receive two sets of output commands (one from each control unit). There are a number of strategies related to software aspects of a fault tolerant controller, including implementation diversity, static and dynamic reconfiguration, that must be considered when implementing a specific application.

The potential design patterns described above are generic, and may be applied to a range of autonomous vehicle behaviors, however, in general the above patterns focus primarily on physical redundancy, and additional capabilities may be needed. For example, fail operational lane keeping might involve storing a few seconds of “look ahead” information in case the camera fails. Thus, a pure focus on physical redundancy may not be fully adequate for implementing autonomous vehicle systems.

Summary

In this paper we described a safety strategy to account for factors relevant to the safety analysis of autonomous systems, namely potential malfunctioning behavior, availability of autonomous driving conditions, and driver attentiveness. Our strategy includes the definition of a controller that monitors and analyzes these factors and helps

ensure fault handling and degradation when one or more of them, necessary for autonomous driving, no longer exist, or will no longer exist at some time in the future. We also provided some potential design strategies for implementing the high integrity hardware and software needed for autonomous driving vehicle systems.

References

1. <http://www.tartanracing.org/>
2. http://vislab.it/Projects/view/32/VisLab's_new_adventure_on_the_Silk_road
3. <http://www.tgdaily.com/hardware-features/51947-google-test-drives-autonomous-vehicles>
4. A. L. Juarez Dominguez. *Feature Interaction Detection in the Automotive Domain*. In the Proceedings of the Doctoral Symposium of the 23rd IEEE/ACM International Conference on Automated Software Engineering (ASE). L'Aquila, Italy. September, 2008.
5. A. L. Juarez Dominguez, J. J. Joyce, and R. Debouk. *Feature Interaction as a Source of Risk in Complex Software-intensive Systems*. In Proceedings of the 25th International System Safety Conference (ISSC 2007), Baltimore, USA. August, 2007.

Biography

Rami Debouk, Ph.D., Staff Researcher, General Motors Company, 30500 Mound Road, Warren, MI 48090-9055, USA, telephone – (313) 820-5358, facsimile – (586) 986-3003, e-mail – rami.debouk@gm.com.

Rami Debouk received a B. Eng. (with distinction) and M. Eng., both in Computer and Communications Engineering, from the American University of Beirut, and a Ph.D. in Electrical Engineering - Systems from The University of Michigan, Ann Arbor, in 2000. Since 2000, he has been with the Electrical and Controls Integration Lab. at General Motors' Research and Development Center, where he is conducting research in the area of safety critical by-wire systems. He has been a consultant to Siemens Corporate Research in the area of failure diagnosis. Other research interests include discrete event systems, decentralized information systems, fault tolerant systems, and software safety. Dr. Debouk is a senior member of IEEE and the System Safety Society, and a member of the Society of Automotive Engineers (SAE) and Sigma Xi. He was named "Engineer of the Year" by the International System Safety Society in 2009.

Jeff J. Joyce, President, Critical Systems Labs, Inc., 618-475 Howe Street, Vancouver, BC, V6C 2B3, CANADA, telephone – (604) 688-2754, facsimile – (604) 628-5692, e-mail – jeff.joyce@cslabs.com.

Dr. Joyce is president and co-founder of Critical Systems Labs, Inc. – an engineering consultancy that specializes in software-intensive, safety-critical systems. He earned a Ph.D. (Computer Science) from Cambridge University (UK) in 1990. Dr. Joyce was the General Chair of the 26th International System Safety Conference. He was named "Engineer of the Year" by the International System Safety Society in 2006.

Barbara Czerny, General Motors Company, 30500 Mound Road, Warren, MI, USA, 48090, Mail Code 480-106-390, telephone – (586) 986-3003, facsimile – (248) 807-3686, email – barbara.czerny@gm.com.

Dr. Barbara Czerny is a Staff Researcher in the Electronic Control Systems Process, Methods, & Tools group at General Motors Research Labs, and has worked in the automotive industry for 24 years. She has over twelve years experience in automotive system safety, and has worked in the areas of system and software safety applied to safety-critical automotive applications such as by-wire systems. She is an ISO technical expert involved in the development of the automotive safety standard ISO 26262. Czerny received a Ph.D. degree in computer science from Michigan State University. She is a member of the System Safety Society.

Joseph D'Ambrosio, General Motors Company, 30500 Mound Road, Warren, MI, USA, 48090, Mail Code 480-106-390, telephone – (586) 986-7291, facsimile – (248) 807-3686, email – joseph.dambrosio@gm.com.

Dr. Joseph D'Ambrosio is a Lab Group Manager for the Electronic Control Systems Process, Methods, & Tools group at General Motors Research Labs, and has worked in the automotive industry for 28 years. He has over 13 years of experience working in the area of safety-critical systems, including steer-by-wire and brake-by-wire applications. D'Ambrosio received a Ph.D. from the University of Michigan, and is a member of the SAE.