Intelligent Vehicle Dependability and Security (IVDS) Project

International Federation for Information Processing (IFIP) Working Group 10.4
on
Dependable Computing and Fault Tolerance

FINAL REPORT

Jaynarayan H Lala
John F Meyer
Carl E Landwehr
Charles B Weinstock
Wilfried F Steiner

1 February 2024

ifip

WG 10.4

Executive Summary

Intelligent vehicles hold great promise for relieving the tedium of driving while simultaneously reducing the tragic injury and death toll on the roads. However, the headlong rush to be the first to market, without adequate safety and security considerations of life-critical control system design, could cause irreparable public harm and set back the promise of autonomous driving.

This concern motivated IFIP Working Group 10.4 (WG 10.4) to initiate a multi-year Intelligent Vehicle Dependability and Security (IVDS) project in June, 2019. WG 10.4 was established by IFIP in 1980 and has contributed significantly to both R&D and application of highly dependable and secure computing systems through workshops, affiliations with conferences, and landmark publications by many of its members.

The vision of the IVDS project is the realization of highly dependable and secure operation of intelligent vehicles, verified and validated with respect to strict dependability (particularly safety) and security requirements by rigorous state-of-the-art methods. Our mission is to engage stakeholders to increase awareness of dependability and safety requirements, promote technical solutions, and provide expert help to governance and regulatory bodies in their rulemaking and oversight roles.

The project has made significant progress towards achievement of its objectives. The purpose of this report is to provide an easily referenced corpus of project results to stakeholders, including autonomous vehicle designers, builders, researchers, and governance and regulatory authorities, and lawmakers.

Representatives from these stakeholders were brought together in two workshops to candidly discuss the current state of practice, the state of the art, vision of desired end state, and the challenges to realizing that vision. Additional engagements with the community included opinion pieces, white papers, academic journal publications, social media presence, and an international panel.

Principal findings of the project point to significant shortfalls in technologies, cost, governance, and societal aspects in achieving the end goal of safe and secure self-driving intelligent vehicles, generally referred to as SAE Level 4 or 5. Self-driving at L4 is subject to limited conditions specified by an Operational Design Domain (ODD). With sufficiently restrictive ODDs, good progress is being made toward achieving this goal. On the other hand, at L5 an intelligent vehicle can drive anywhere under any conditions. Widespread deployment of such vehicles on public roads, much hyped at the outset of the project, appears to have receded into the more distant future.

There are some limited-scope success stories. For example, in highly constrained environments, goods are being transported by driverless trucks, monitored remotely, on fixed, preplanned routes under restricted speed and weather conditions. Passengers are also being provided autonomous vehicle rides in a few cities under similarly controlled conditions.

However, progressing to L5 autonomy will require orders of magnitude improvements in machine learning algorithms' capabilities to classify objects correctly, to keep determined cyber attackers at bay, and to create affordable vehicle guidance, navigation, and control systems that are fail operational / fail safe, the minimum fault-tolerance requirement for L5 driving. Quantitative

standards of safety must also be developed, to complement the qualitative ISO/PAS 21448: Safety Of The Intended Functionality (SOTIF) and other standards. These would provide the community a common measurement yardstick: a specific design criterion for the industry, an objective standard to certify vehicles for the regulators, and a specific goal for technology developers and the research community to aim for. Assurance and certification of systems incorporating new, nondeterministic technologies such as machine learning (ML), also demands innovative approaches. Many, though not all, of these challenges have been faced by other industries in the past, the most prominent being the aerospace and defense sector. The automotive community should apply lessons from their experience to dealing with a highly complex road environment.

While the challenges cannot be minimized, history of innovation shows that it is feasible to overcome them and to achieve the vision of autonomous road transport that can potentially save millions of lives and avoid many more needless injuries. But it will require the autonomous vehicle industry to make dependability and security top priorities and all stakeholders to collaborate and coordinate their efforts with a single-minded purpose and end state. If that does not happen, we would lose the only opportunity we have had since the inception of automobile transport to improve its safety by many orders of magnitude.

1. Problem Statement

Driving a road vehicle is a complicated endeavor requiring a seamless interaction of multiple cognitive abilities. These include attention, perceptuomotor skills, memory, and decision-making. Automating these abilities within an intelligent vehicle (IV) poses challenges which if met could provide greater road safety, reduced road congestion, increased productivity, and more independence for older adults, especially those with disabilities. In particular, a major benefit would be reducing tragic injuries and deaths due to impaired human drivers, who currently cause approximately one-third of road fatalities.

Among the many problems that need to be addressed in this regard, we chose to focus on the realization of highly dependable and secure operation of such vehicles, verified and validated with respect to strict dependability (particularly safety) and security requirements by rigorous state-of-the-art methods. The types of vehicle intelligence assumed are principally those associated with SAE Driving Automation Levels 3-5, per SAE Int'l Recommended Practice J3016.

Although an IV can (and should) be assisted by an infrastructure aimed at enhancing its dependable and secure operation, for the most part, the problem addressed by the IVDS project concerns what the road vehicle, per se, must accomplish when either assisting or replacing a human driver. Accordingly, factors external to a single IV such as road type, time of day, other vehicles, pedestrians, road signs, traffic volumes, weather, and visibility conditions, etc. are regarded as the vehicle's environment. Although the nature of this environment is a critical concern (e.g., how safety of an IV varies with ODD), how it might be altered to improve IV dependability and security lies outside of the project's scope.

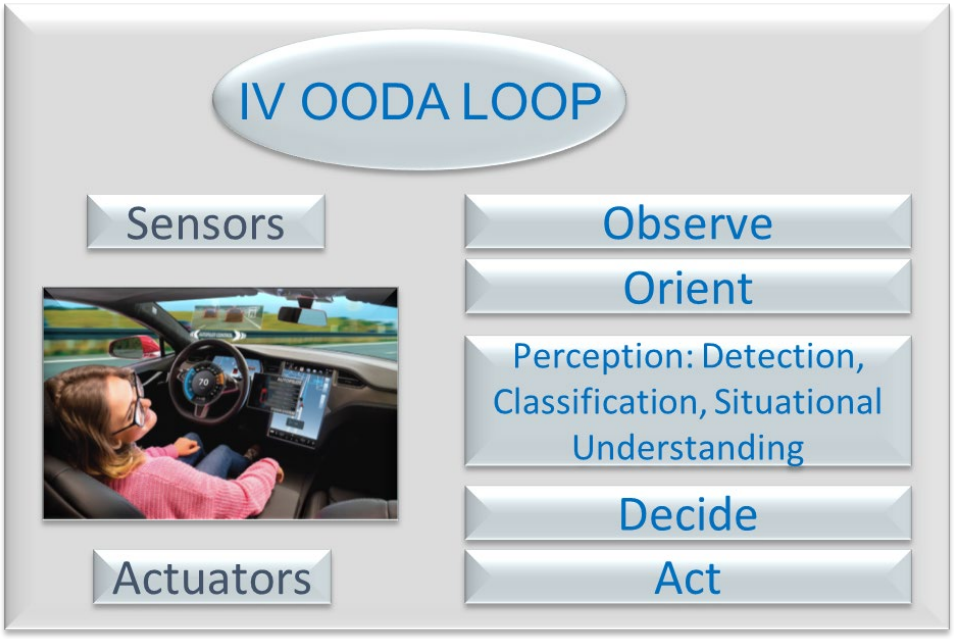

Figure 1.1 Intelligent Vehicle OODA Loop

Figure 1.1 shows the major functions that an IV must master. These are organized as components of the OODA Loop: Observe, Orient, Decide, and Act. Some of the functions have been

implemented quite successfully, with a very high degree of dependability. These include sensing own vehicle state such as speed, acceleration, and engine parameters; navigation related state such as position, and direction of travel on a map; and actuator state such as steering, braking, and engine power settings. To replicate overall situational awareness of a human driver, additional observations and correlation of that data into orientation are necessary. Some of these functions are easier than others to realize and have also been in widespread use for a while. Adaptive cruise control, for example, measures distance to the vehicle in front as well as its velocity and acceleration, to automatically manage own speed to keep a safe distance. Similarly, presence of vehicles in adjoining lanes is monitored to warn of unsafe lane changes. The more challenging parts of creating a human-like situational awareness include: 1) detection of all objects in all weather conditions, 2) classification of objects (e.g., a pedestrian, stroller, bicycle, debris, emergency vehicle, train, road sign, etc.), 3) detection of stopped vehicles, and 4) lane classifications (one way, two ways, wrong way, etc.) even when there are no clear lane markings. Additionally, human drivers are also very adept at forming a mental picture of an unfolding situation over a longer time period such as might be caused by abnormal and unexpected behavior of other actors in the environment – something very challenging for an IV.

Since intelligence in an IV augments natural intelligence of a human driver, it is no surprise that artificial intelligence (AI) technologies such as machine learning (ML) are being employed to create real-time situational awareness. This results in an added degree of difficulty since attributes of dependability are based on an underlying concept of “failure” (a transition from correct to incorrect service delivery). Generally, characterizing failure of an AI-enabled system (whether or not autonomous) is difficult, since what’s meant by correct (desired) or incorrect (undesired) service is elusive. In particular, this issue calls for innovative techniques throughout the IV development cycle, ranging from architecture specification through IV dependability and security assurance.

2. Background

IVDS Project Genesis

During the June 2018 meeting of IFIP Working Group 10.4, the need for more active member engagement in the form of one or more dedicated projects was suggested by John Meyer. In turn, Carl Landwehr volunteered to chair a committee to solicit project proposals, evaluate submissions, synthesize them where needed, and present the most promising candidates to the membership. Out of nine submissions, three emerged from this process and were submitted for approval at the June 2019 WG meeting. The IVDS proposal[1] was one of two approved, with Jay Lala volunteering to lead the project.

Although this was the first WG 10.4 project with an IV focus, contributions by WG members regarding the dependability (particularly safety) and security of intelligent autonomous road vehicles began well before the project's inception and continued throughout the project's duration. In addition to journal articles in this regard, WG members have presented papers/talks on the subject at numerous related conferences/workshops over the past decade. These include (but are not limited to):

- IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)
- DSN workshops on Safety and Security of Intelligent Vehicles (SSIV);10 since 2014
- DSN workshops on Dependable and Secure Machine Learning (DSML); 6 since 2018
- IFIP WG 10.4 Workshops; 5 since 2015, including two IVDS workshops
- International Conference on Computer Safety, Reliability and Security (SAFECOMP)
- European Dependable Computing Conference (EDCC)

Regarding the choice of *Intelligent vehicle* (IV) in the project's title, there are a number of other terms that refer to a road vehicle (car, bus, truck, etc.) or a collection thereof, wherein the role of a human driver is either reduced or eliminated. These include *autonomous vehicle* (AV), *connected autonomous vehicles* (CAVs*), intelligent transportation system* (ITS), *self-driving car,* and *automated vehicle*. The choice of IV was motivated primarily by considerations delineated in the problem statement of Section 1. In particular, the term "intelligent" was chosen to emphasize that replacing the intelligence required of a human driver (see the OODA Loop of Fig. 1.1) necessitates autonomy that is likewise (artificially) intelligent. The use of "vehicle" was chosen to indicate that the scope of the project was primarily limited to a road vehicle, per se. This is in contrast to terms such as CAV, where C refers to V2V or V2X connectivity, or ITS where S can encompass an entire intelligent vehicle ecosystem.

An exception to our "vehicle" focus appears in the statement of Goal 4 (see below); however, this extension was not addressed during the course of work on the project. We should also note here that in this project, we use industry standard terminology such as self-driving SAE levels and do not use company-specific terminology such as Full Self Driving (FSD).

[1] https://www.dependability.org/wg10.4/ivdswiki/images/8/8f/Intelligent_Autonomous_Systems_Project.pdf

Project Vision, Mission, and Goals

As conceived at the outset of the project, its vision, mission, and goals were stated as follows in a white paper[2] distributed at *The Autonomous* in Vienna, AT, Sept 2019.

Vision

The vision of the project is the realization of highly dependable and secure operation of intelligent vehicles (IVs), verified and validated with respect to strict dependability (particularly safety) and security requirements by rigorous state-of-the-art methods. Types of vehicle intelligence assumed in the vision are mainly those associated with SAE levels 3-5 (ADS) (per SAE Recommended Practice J3016, V 0620180), although level 2 (ADAS) may also be addressed with respect to certain issues.

Mission

In pursuit of this vision, the project's mission is to facilitate awareness and provide pro bono counsel to both automotive stakeholders and relevant standardization/regulatory bodies. Avenues for such provision include the invitation of researchers and government/industry stakeholders to workshops regarding specific IVDS issues. These will be organized either by WG 10.4 or jointly with other professional groups having similar interests, e.g., WGs in IFIP TCs such as TC 11 (Security) and TC-12 (AI), and IEEE TCs and ACM SIGs such as the *TC on Self Driving Cars* (IEEE Intelligent Transportation Systems Society), *TC on Autonomous Ground Vehicles and Intelligent Transportation Systems* (IEEE Robotics and Automation Society), and the *SIG on Artificial Intelligence* (ACM). Another important avenue is meeting with industry and government in their own environments, e.g., IV-related conferences and workshops organized by industrial groups and governmental agencies.

Goals

In keeping with the mission, the goals of the project correspond to acquiring and disseminating knowledge with respect to a number of important concerns regarding dependable and secure IVs.

Goal 1: Measures

Several interesting questions deserve consideration in this regard. Generally, what kind of measures are appropriate for assessing dependability and security during various phases of IV development and deployment? Are there dependability measures other than safety that need to be addressed in this context? What are the most useful measures of IV security? Can security of an IV be measured quantitatively? A number of publications (journal papers, conference papers, white papers and technical reports by government and non-profit organizations, etc.) have been published in this regard. However, most focus on safety and tend to reflect legacy thinking relating to safe SAE level-0 operation of vehicles.

[2] https://www.dependability.org/wg10.4/ivds/IFIP_WG10.4_White_Paper_on_IVDS-V2.pdf

Goal 2: Standards

Many standards for IV operation at various levels of automation already exist. These requirements are authored by various international bodies and national organizations in the Americas, Europe, and Asia. (Although not advertised as a standard, SAE J3016 is highly suggestive of what needs to be standardized.) These documents deserve to be studied carefully with regard to dependability and security requirements. In particular, are they consistent? Do they employ measures of the type addressed in Goal 1? If so, are there stated bounds on measure values that need to be adhered to?

What's likely needed in this regard are dependability specifications of the type that have been applied to aircraft and spacecraft computing systems. For example, if failure is identified with a crash causing fatalities, the allowable failure rate (e.g., fatal crashes per hour) is no greater than a very small number such as 10^{-8}.

Goal 3: Employment of Machine Learning Algorithms

If a vehicle is intelligent then, de facto, it is artificially so, i.e., AI is employed, particularly the use of machine learning (ML) techniques. However, there are general concerns about the use of statistical ML algorithms in safety-critical applications and, hence, their use for IV control. Should employment of such algorithms therefore be avoided in IVs? If not, are there methods of ensuring that unpredictable anomalies in outcomes are detected and their adverse effects on the IV are mitigated?

Goal 4: Key Design Methods

To satisfy strict dependability and security requirements, key design methods will need to be discussed with automotive stakeholders. These will certainly include techniques for tolerating both accidental and intentional faults that are familiar to the dependable computing community. Also included are techniques more specific to IV safety, e.g., those specified in ISO 26262. However, with regard to the concerns noted in Goal 3, new methods may also be needed to tolerate the vagaries of non-determinism.

Since safe and secure operation of IVs transcends that of a single vehicle, system-of-systems behavior also needs to be considered, along with a hierarchy of related policies. In particular, attention should be given to highway flow control, the interaction of multiple vehicles and the potential for cascading error scenarios.

For example, to what degree does design diversity help or hinder achievement of goals? Could a common mode error or security flaw create gridlock or catastrophic failure? Are multiple discipline engineering solutions thoroughly integrated and analyzed? What is the potential for mode confusion and emergent behavior?

Goal 5: Assurance and Certification

IVs present some interesting problems with regard to assuring and certifying that dependability and security requirements are satisfied. These problems cut across all the usual means of verifying and validating safety-critical and security-critical systems. They are due, in part, to the fact that IVs are intelligent and, therefore, less predictable than more usual deterministic systems. Indeed, what needs to be accomplished in this regard depends in no small extent on first resolving

concerns expressed in Goals 1-4. Finally, there's the question as to who is going to be responsible for such assurance and certification. Roles and responsibilities of regulatory, oversight and governance bodies need to be clearly articulated.

IVDS Project Team

During the 4.5-year duration of the project, the following individuals made significant contributions to the project's success.

Dr. Jaynarayan Lala (Project Lead), Raytheon, an RTX Business, US jay.lala@rtx.com
Prof. John Meyer, University of Michigan, US jfm@umich.edu
Prof. Carl Landwehr, University of Michigan, US carl.landwehr@gmail.com
Dr. Charles Weinstock, Carnegie Mellon University, US weinstock@conjelco.com
Prof. Homa Alemzadeh, University of Virginia, US ha4d@virginia.edu
Dr. Wilfried Steiner, TTTech, AT wilfried.steiner@tttech.com
Prof. Cristina Nita-Rotaru, Northeastern University, US cnitarot@gmail.com

3. IVDS Principal Contributions

The principal contributions of the project are found in the results of two IVDS workshops, an IFIP Diamond Jubilee panel, two papers to raise public awareness, and external activities stimulated by project participants and activities. These contributions are summarized in this section.

***First Workshop**, January 29 - February 1, 2021 (virtual).* The workshop results, including slides and video of the presentations, are available to the public on the web at https://ivds.dependability.org/ivds2021/.

Organization of the workshop commenced in the summer of 2020. Planning an in-person workshop was infeasible because of the raging Covid pandemic. Workshop development was coordinated through regular bi-weekly Zoom sessions of the project committee members.

The workshop goals were to outline key challenges facing the realization of highly dependable and secure operations of intelligent vehicles; have subject matter experts address some of these challenges and begin a dialog with some of the IVDS stakeholders, such as automotive industry, government and other standardization/regulatory bodies. In addition, the workshop aimed to identify specific actions, both short term and long term, to achieve the IVDS project vision, mission and goals. As a way to motivate presentations and discussions, a hypothesis was proposed: *Level 3 autonomous vehicles cannot be made acceptably safe with current technology and practices*. Presenters were asked to respond to this hypothesis in the context of their sessions.

To maximize participation from individuals located literally around the world, the workshop was convened as three 2.5-hour sessions held Friday, January 29 – Sunday, January 31 from 9am to 11:30am ET each day. The first day focused on interaction between humans and semi-autonomous systems. The following day explored autonomous vehicle perspectives from industry, and the third day discussed verification and validation concerns for intelligent vehicles. An optional session to facilitate further discussion among the participants was arranged for Monday, February 1, and was well attended.

The workshop attracted more than 70 participants and generated many favorable comments. With respect to the hypothesis, the consensus of the participants was that indeed L3 autonomous vehicles cannot be made acceptably safe with current technology and practices.

***Second Workshop**, June 24-25, 2022 (Alexandria VA) (in person).* The workshop program and speakers can be found at https://ivds.dependability.org/ivds2022/.

The focus of this workshop was to explore how to survive cyber attacks on safety-critical functions of intelligent vehicles. A workshop goal was to discuss design solutions, quantitative cyber-survivability measures, and verification and validation with regard to their impact on AV safety. Planning for this first in-person workshop since the onset of the Covid-19 pandemic began in fall, 2022. The organizers were very grateful to Chris Walter for his substantial and successful efforts in locating both a workshop venue in Alexandria, Virginia, and arranging an excursion to the Udvar-Hazy museum. The workshop attracted nearly 60 participants.

Multiple stakeholders came together to share their diverse perspectives on the main theme of the workshop. Industry representatives described their current approach to detection and mitigation of cyber threats that may impact vehicle safety. They also pointed out the cost constraints of what methods and techniques can be used while keeping their businesses profitable. Researchers

from academic institutions and industry shared novel ideas for countering cyber threats. The workshop audience also heard a regulatory viewpoint from a government representative. And, last but not least, legal implications were explored by a professor of law. All participants left with an increased awareness of, and appreciation for, the challenges and constraints of other stakeholders.

***IFIP 60th Anniversary Panel**, October 18, 2021 (virtual).*

On behalf of WG 10.4 and in response to IFIP's call for proposals to participate in its "60th Anniversary Future of Information Processing Series," the project proposed a panel entitled "Autonomous Vehicle Safety and Security: An Information Processing Imperative." It was subsequently selected by IFIP (June 2021) as one in a 10-panel series to take place virtually during the last half of the year https://ifip.org/jubilee60/?r=events.

As the 6th event in the series https://ifip.org/jubilee60/?r=event6, the panel was moderated by Prof. John Meyer of the University of Michigan and featured four distinguished panelists: Prof. Ravi Iyer of the University of Illinois at Urbana-Champaign, Prof. Missy Cummings, then of Duke University, Dr. Wilfried Steiner of TTTech, and Prof. Philip Koopman of Carnegie Mellon University. The event benefitted significantly from the diversity of the panelists' expertise with various aspects of road vehicle autonomy. Moreover, all shared a common interest in the paramount issue of intelligent vehicle safety. During the session, the opinions they expressed regarding both existing problems and potential solutions for safe IVs signaled "not so fast" to early deployment of this technology.

More than 150 participants attended this event. Based on questions asked, they represented 17 time zones ranging from the west coast of North America to Japan. 50+ questions were posed throughout the session, where during a limited Q&A period, the panelists were able to address 10 verbally and another 20 in writing. The audience also participated in a poll concerning use of safety-related functions of currently deployed Automated Driving Assistance Systems (ADAS).

Public reports drafted and published to raise public awareness

To bring the concerns about the dependability and security of intelligent vehicles to the broader technical community, project members drafted and published a "Viewpoint" article in the *Communications of the ACM*. Drafting of the piece was initiated in the fall of 2019 and it was published in the September 2020, issue of the magazine, with the title "Autonomous Vehicle Safety: Lessons from Aviation."[3]

The project also sought to reach the general public and in particular those responsible for legislation and regulation in the U.S. In the spring of 2021, project members drafted an opinion piece that appeared in *The Hill,* July 21, 2021 under the title "'Move fast and break things' won't work for autonomous vehicles."[4]

[3] https://www.dependability.org/wg10.4/ivdswiki/images/6/67/CACM_Viewpoint_09-2020.pdf

[4] https://thehill.com/opinion/technology/563915-move-fast-and-break-things-wont-work-for-autonomous-vehicles/h

External activities

In addition to workshops and publications, the project engaged in other outside activities to stimulate awareness of the issues in intelligent vehicle dependability and security. These included:

Project leader Dr. Jay Lala presented “Autonomous Vehicles: Safety Measures and Benchmarks for Perception & Cognition Functions,” at the Association for the Advancement of Artificial Intelligence (AAAI) Fall Symposium Series on Nov. 5, 2021. The symposium details can be found at https://aaai.org/conference/fall-symposia/fss21/.

Project members joined with other computing professionals starting in fall, 2022, to participate in the Association for Computer Machinery (ACM) U.S. Technical Policy to craft an ACM position on the adoption of autonomous vehicle technology.

Archive of relevant materials created for general use

An archive of materials relevant to the project’s focus was created and is available to the public. Although entries are no longer being added to it, it includes project documents as well as many links to publications relevant to the ongoing development and deployment of intelligent vehicles. It may be found at https://ivds.dependability.org/ivds-additional-material.html.

4. Current Situation

At the time of this writing (2023) dependable and secure intelligent vehicles have not yet entered the mainstream market. Some of the remaining key issues are as follows.

Commonly accepted precise qualitative and quantitative safety/security goals of self-driving cars must be developed and societally agreed upon. Said goals may differ concerning Operational Design Domains (ODDs). Regulations like the European Regulation 2019/2144 on modern vehicle technology[5]

> *"The manufacturer shall define the acceptance criteria from which the validation targets of the ADS are derived to evaluate the residual risk taking into account existing accident data,* data on performances from competent and carefully driven manual vehicles and technology state-of-the-art.*
>
> **For instance based on current accident data on buses, coaches trucks and cars, an indicative aggregated acceptance criteria of 10^{-7} fatalities per hour of operation could be considered for market introduction of ADSs for comparable transport services and situations. The manufacturer may use other metrics and method provided it can demonstrate that it leads to an equivalent level of safety."*

may serve as a basis but must be further refined and extended.

Several standards have been newly developed or extended to reflect the specific needs of self-driving cars. Examples are *ISO 26262 Road vehicles Functional safety, ISO/PAS 21448 Safety of the Intended Functionality, UL 4600 Standard for Safety for the Evaluation of Autonomous Products, ISO/SAE 21434 Road vehicles - Cybersecurity engineering*. While these and other standardization activities advance the global state-of-the-art in the safety and security of self-driving cars, they must evolve as self-driving cars get closer to series production. For example, the second IVDS workshop, themed "How to Survive Cyber Attacks on Safety-Critical Functions of Intelligent Vehicles," discussed potential improvements to the *ISO/SAE 21434 Road vehicles - Cybersecurity engineering* standard.

Although commonly accepted safety and security goals are missing, it is our expert opinion that said goals are so demanding that they require self-driving cars to implement fault-tolerant distributed computer systems. Architecting such computer systems is challenging as it requires an effective system decomposition into nearly independent subsystems that form fault-containment units. Some architectures have been proposed, for example, an architecture by Kopetz[6]. Further architectures are analyzed in The Autonomous [7].

---

[5] European Commission. (2022). Appendix 1, 7.1.1. Rules for the application of Regulation (EU) 2019/2144 of the European Parliament and of the Council as regards uniform procedures and technical specifications for the type-approval of the automated driving system (ADS) of fully automated motor vehicles. https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12152-Automated-cars-technical-specifications_en

[6] Kopetz, H. (2022). An Architecture for Safe Driving Automation. In Principles of Systems Design: Essays Dedicated to Thomas A. Henzinger on the Occasion of His 60th Birthday (pp. 61-84). Cham: Springer Nature Switzerland.

[7] The Autonomous. (2023). Safe Automated Driving: Requirements and Architectures. https://www.the-autonomous.com/wp-content/uploads/2023/12/wg-safetyarchitecture-full-report-a4.pdf

Assuring a self-driving car meets its safety and security goals is another challenge. It appears unlikely that such assurance can be achieved by test driving only because of the huge number of driving miles required[8]. Effective combinations of validation procedures need to be further explored.

Self-driving cars have already been involved in accidents, some lethal. In several cases, the share of liability between car manufacturers and car operators appears unclear. We expect conclusions and rulings from ongoing and future court proceedings will influence legislation and result in technical requirements for self-driving cars.

While the state-of-the-art is currently only forming, some automotive companies have already started to roll-out self-driving technologies for different use cases with varying degrees of success and failure. Companies offering self-driving technologies in complex ODDs still require a human supervisor. However, there are numerous registered incidents in which the human failed this supervision task. Some of these incidents were lethal. Some further companies offer robo-taxi services. These services are typically restricted to a few cities with simple ODDs. Even in these relatively simple use cases, incidents have already been reported. In another use case, driverless trucks operate in some highly restricted ODDs. To our knowledge, incidents have not yet been reported in this use case.

[8] Paddock, S. M., & Kalra, N. (2016). Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability?

5. Conclusions & Path Forward

The IVDS project vision has been the realization of highly dependable and secure operation of intelligent vehicles, verified and validated with respect to strict dependability (particularly safety) and security requirements by rigorous state-of-the-art methods.

Principal findings of the project, conducted over the past four plus years, point to significant shortfalls in technologies, cost, governance, and societal aspects in achieving the end goal of safe and secure SAE Level 4 or 5 self-driving intelligent vehicles.

The consensus of nearly 70 participants at the first IVDS workshop, including representatives from industry and academia, was that at that time, even Level 3 autonomous vehicles could not be made acceptably safe with current technology and practices. However, with sufficiently restrictive ODDs, good progress is being made.

Recent developments support these findings. The driverless taxi permit for a major autonomous vehicle manufacturer and operator, General Motors' Cruise Division, was cancelled by the state of California due to an unfortunate accident and subsequently the operator withdrew its robo-taxis from all US cities[9].

The US National Highway Traffic Safety Administration (NHTSA) recently issued a recall of more than 2 million Tesla vehicles to ensure drivers are paying adequate attention when "autopilot driver assistance system" is activated[10].

A McKinsey & Co survey of 86 stakeholders reveals that "2023 was a tipping point for the autonomous-vehicle industry"[11].

Top critical technology challenges are prediction & decision making and perception software. Leading players also observed "….significant setbacks, stopped or reduced their operations, or exited the market entirely." The survey also concluded that "The timeline for autonomous-vehicle is extending." L4 robo-taxis are expected to become commercially available at a large scale by 2030.

In pursuit of the IVDS project's vision, its mission was to facilitate awareness and provide pro bono counsel to both automotive stakeholders and relevant standardization/regulatory bodies.

The project largely accomplished the mission through its workshops, an international panel, publications and presentations in scientific journals, and even an opinion piece in a Washington, DC newspaper. Legal viewpoints were also considered in these discussions. Regrettably, the regulatory and governance representatives were mostly absent.

All the project outputs, including workshop presentations and video recordings by leading researchers and practitioners in the field, have been organized into an easily accessible website ivds.dependability.org. The corpus of knowledge also includes many relevant safety standards. The reader is encouraged to explore these artifacts for a deeper dive.

---

[9] https://www.npr.org/2023/12/30/1222083720/driverless-cars-gm-cruise-waymo-san-francisco-accidents

[10] https://static.nhtsa.gov/odi/rcl/2023/RCAK-23V838-3395.pdf

[11] https://www.mckinsey.com/features/mckinsey-center-for-future-mobility/our-insights/autonomous-vehicles-moving-forward-perspectives-from-industry-leaders?utm_medium=DSMN8&utm_source=LinkedIn&utm_user=14419234616064771

Going forward, it is evident that the paradigm of “moving fast and breaking things” is not the right approach to replace human-driven vehicles with autonomous vehicles, just as it wasn’t for aviation at the dawn of the jet age. Had it been, with no quantitative safety standards, no Federal Aviation Administration or its global counterparts to oversee safety of airliners, airplane manufacturers and airlines would have been allowed to offer commercial air services without certification. They might have adopted “move fast and learn”, resulting in unnecessary crashes and loss of innocent passenger life. We would never have reached the state today where 10+ million flights transport over 700 million passengers yearly in the US with an exemplary safety record of no fatalities in the last 15 years.

A path exists to realize the end-state of road transport that is just as safe as aviation, if we learn the lessons from aviation community in how to design, build, test, validate, operate and regulate safety-critical systems.

Acknowledgements

The IVDS team wishes to express its appreciation of the support provided by WG 10.4’s leadership and members throughout the duration of the project. Our sincerest thanks also to the speakers and panelists of the two IVDS workshops and the IFIP 60th Anniversary panel, and to all those who attended these events or otherwise contributed behind-the-scenes to the project's success.