DOI:10.1145/3411053

Jaynarayan H. Lala, Carl E. Landwehr, and John F. Meyer

▸ **Terry Benzel,** Column Editor

Security

Autonomous Vehicle Safety: Lessons from Aviation

How more than 25 years of experience with aviation safety-critical systems can be applied to autonomous vehicle systems.

AUTONOMOUS VEHICLES SEEM to hold great promise for relieving humans of the boring task of guiding a car through congested traffic or along a monotonous turnpike, while at the same time reducing the annual highway death toll. However, the headlong rush to be the first to market, without adequate considerations of life-critical control system design, could cause irreparable public harm and ultimately set back the promise of autonomous driving. With the current goal of being at least as safe as human driving, espoused by business leaders as well as some regulatory agencies, the annual death toll attributed to automation killing innocent people, just in the U.S., would be approximately 36,500 per year or 100 per day. Think about that for a minute!

This column highlights this important dependability need, the dire consequences of falling short, and how leveraging the knowledge gained by the aviation industry in operating safety-critical flight control systems without fatalities for over a quarter century can help avoid this outcome.

SAE International has defined six levels of autonomy for on-road motor vehicles in SAE J3016, where an updated graphic summary of the levels was released last year.[10] An excerpt of this visual describing "What does a human in the driver's seat have to do?" is displayed in the figure on p. 29.

Numerous automobile companies[2] are racing to be the first to market. Ford says it will have an L5 vehicle in operation by 2021. More ambitiously, Toyota announced in February 2019 that it plans to have a self-driving vehicle ("the most intelligent supercomputer on wheels") available for purchase within a year. To achieve this objective, Toyota's vice president in charge of software says "Our goal is to teach a Silicon Valley mindset here."

Waymo, a spinoff of Google's self-driving car project, is already rolling out a fully autonomous ride-hailing service in Phoenix, AZ, and has run road tests of autonomous "big rig" trucks in Atlanta, GA. Volkswagen announced it would field a fleet of self-driving vehicles by 2021, and has recently teamed with Ford. As early as 2016, Tesla announced all cars it produces have the hardware needed for L5 driving capability; evidently it is just a small matter of programming.

Every manufacturer has a safety and security argument to go along with its autonomous vehicle control system (AVCS) plans. For cybersecurity, they may point to guidelines published by the U.S. Department of Transportation[12] or best practices published by the Auto-ISAC.[1]

Twelve leading companies have collaborated on a 150-page white paper "Safety First for Automated Driving,"

IMAGE BY SCHARFSINN

working toward "industry-wide standardization of automated driving,"[11] which builds on guidelines and standards worldwide, including some still under development.[4] It describes both development processes aimed at achieving "safety by design" and verification and validation of elements and systems at L3 and L4 autonomy. Use of Deep Neural Nets to implement safety-related elements is addressed in an appendix in the white paper.

This effort appears comprehensive and a strong step in the right direction, but, as its authors recognize, it is a work in progress. Further, as observers have noted,[3] it is strictly a voluntary initiative among a set of companies, outside of the usual channels for standards development. Moreover, the effort is generally focused on standards companies apply to their internal design and development processes without external review or certification. They do not provide a quantifiable level of safety or security performance in terms of, for example, expected failure rate of control systems per hour or mile of operation.

While the quest for autonomous vehicles may seem novel, we have been here before. Aviation provides a successful model with many parallel challenges, including hard real-time control, dependability commensurate with the adverse economic and life-threatening consequences of failures, scalability, affordability, and non-technical factors such as certification and governance. A case study of the most relevant avionics application—full authority, full time, fly-by-wire (FBW) aircraft flight control systems (FCS) is instructive.

The dependability requirements, including safety and reliability, were defined only after several years of discussions with NASA and the FAA in the mid-1970s. NASA was sponsoring research in FBW FCS to enable commercial aircraft to be more fuel efficient by taking advantage of statically unstable aircraft designs. An aircraft of this design would require computer control to maintain its stability.

The FAA's initial proposal was to mandate FCS to be as reliable as the wings of the airplane, that is, the FCS should never fail. The rationale was that wings never fall off and the FCS is just as integral to an aircraft's safety as its structure. However, from an engineering viewpoint, that was not a requirement that one could design to.

SAE levels L0–L5.

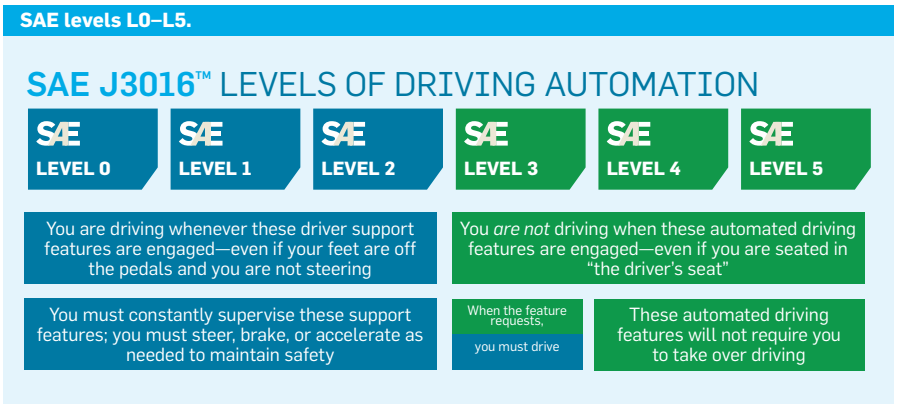

After further discussions, the FAA defined a quantitative requirement in a one-page memo to NASA, in May 1974. Key excerpts of the memo are quoted below.

"Section 25.1309 of the Federal Aviation Regulations requires that airplane systems be designed so that occurrence of any failure conditions (combinations of failures in addition to single failure considerations) which would prevent the continued safe flight and landing of the airplane is extremely improbable. The FAA has accepted substantiating data for compliance with that requirement which shows by analysis that the predicted probability of occurrence of each such failure condition is 10^{-9} per hour of flight."

"We further believe that failure of all channels on the same flight in a "fly-by-wire" flight control system should be extremely improbable."

In the late 1970s, Draper Lab and SRI International produced two competitive designs, Fault Tolerant Multi-Processor (FTMP) and Software-Implemented Fault Tolerance (SIFT), respectively, under contract with NASA Langley Research Center. These designs were realized in flight-worthy computers by Collins Avionics and Bendix, respectively, and subjected to many theoretical and experimental tests. The architectures relied on hardware and process redundancy; real-time fault detection, identification and reconfiguration; and software fault-tolerance, among many other dependability-enhancing techniques.

While the quest for autonomous vehicles may seem novel, we have been here before.

Verification and validation techniques included proofs of correctness, hardware and software fault injections, measurement of computer response to such events, and Markov reliability models with some of the model parameter values determined via experimentation.

These pioneering research and development efforts resulted in fundamental architectures, designs, theories, and certification methods that continue to shape today's FBW systems.[5]

What Can We Learn from the Aviation Example With Respect to Autonomous Vehicle Dependability Requirements?

First, vehicle control imposes hard real-time requirements, and stringent low latency, just as FCSs do. A control or communication failure during critical vehicle maneuvers can lead to a cascading series of life-threatening accidents. The autonomous control system must detect any consequential fault and take corrective action within fractions of a second to keep the vehicle under control.

Although aircraft FBW systems must function for the duration of the flight, ground vehicles have the luxury of pulling off the road in case of a malfunction.

Still, the control system must, at a minimum, continue to function correctly for the time it takes to maneuver the vehicle to a safe place while also configuring itself into a safe state.

The control system must continue to function correctly after a fault, that is, it must be designed so there are no single points of failure. In case of faults or errors, the system must compensate, and still produce correct results in a timely manner long enough to reach a safe place and configure the vehicle into a safe state. A graceful degradation to a limited functionality Fail-Operational requirement would seem to be adequate.

Aircraft FBW systems use masking, redundancy, and sophisticated reconfigurations to continue to provide full functionality after a fault. Autonomous vehicle control systems can be simpler and less expensive by providing a limited Fail-Operational architecture. Some inspiration can be gained from early aviation experience.

First-generation jumbo-jets, in the early 1970s, used computers to provide "all-weather" auto-land capabilities for Cat IIIB conditions: zero visibility, zero ceiling. They had safety and reliability requirements and mission times very similar to those for autonomous vehicles. There could be no single point of failure and the system had to be operational for only several minutes: the duration of approach and landing. The architectures ranged from dual redundant self-checking pair of computers (Lockheed TriStar L-1011), to duplex channels, each with dual fail-disconnect computers for pitch, roll, and yaw axes (Douglas DC-10) to triple redundant analog computers (Boeing 747).

How Does AVCS Reliability Correlate With the FAA's Mandated Failure Rate of 10^{-9} per Hour for FCS?

Table 1 summarizes recent U.S. motor vehicle death rates, which translate to a fatality rate of 5×10^{-7}/hour per vehicle, assuming an overall average speed of 40 MPH. What would be an acceptable failure rate of an L5 control system? Table 2 illustrates the effects of some alternative rates.

Many people in the industry say just slightly better than status quo would save lives. But that would mean nearly 100 people being killed by road vehicles every single day. Would society accept such mayhem attributed to machines? Recall the public reaction when a single pedestrian was killed by an autonomous Uber in Tempe, AZ.[13]

The FAA set the failure rate for FBW systems at two to three orders of magnitude *smaller* than that of human pilots who are highly trained for safety. Shouldn't society set similar goals for autonomous vehicle safety? Even a failure rate of 100 times better than an average driver, whose safety behavior is unlikely to approach that of a pilot, would still result in 365 U.S. deaths per year attributable to AVCSs. By comparison, no single death has been caused by FBW systems in over a quarter century of operations worldwide.

This is a once-in-a-century opportunity to ride on the revolutionary remaking of ground transportation to make it as safe as aviation, a gift to humanity worldwide.

Regarding safety standards for road vehicles, ISO 26262 concerns the functional safety of on-vehicle electrical and electronic (E/E) systems, where components are ranked according to an ASIL (Automotive Safety Integrity Level), the most critical being level D. However, it should be noted that ISO 26262 (and more generally the auto industry) shies away from quantitative safety requirements, particularly with regard to fatalities.

While drafting ISO 26262, the automotive industry recognized that, in addition to faults in E/E equipment, unsafe behavior could likewise be caused by faults in the specified functionality, spinning off a complementary functional safety standard ISO/PAS 21448 or SOTIF (Safety of The Intended Functionality). The latter presumes the realized system is fault-free and focuses on reducing safety risks due to insufficiencies in the intended behavior. However, these are process-related requirements and not quantitative. Furthermore, the automotive industry "self-certifies" compliance to these standards. Self-certification of the Boeing 737 MAX led to the MCAS system, at the center of the two crashes, being declared non-safety-critical.[7]

Additionally, a safety-assurance case must go beyond the simplistic, non-scientific miles driven and fatal accidents recorded for autonomous vehicles. The data must include the envelope of corner/edge cases explored, realism of testing conditions, unplanned disengagements, and incorrect decision making. Further, this effort should be complemented with analytical models, fault injections, and proofs-of-correctness, where appropriate.

What Can We Learn from Aviation Human Factors?

Much of the automotive industry is moving successively from L0 (fully manual) to L5 (fully autonomous), believing this step-wise increase in autonomy is the safest way to proceed. Counterintuitively, this approach poses a real human-factors challenge. L3 will be a semi-autonomous mode where routine driving is performed by the control system and the human driver's role will be to intervene in emergencies/malfunctions.

In the cockpit, highly trained pilots are primed to recognize unusual situations quickly and take corrective action. Recurring simulator training focuses on dealing with emergencies. Cockpits and flight control systems are designed to optimize human-machine interaction in the safest way possible.

Ordinary driver's license certification requires no such training, nor is it reasonable to expect the public at large to perform at this level. The rarity of emergencies requiring human intervention makes it nearly impossible to keep the driver, who is busy doing other things, sufficiently engaged to take over within a fraction of a second of an alert.

The proper balance between fully automated functions and relying on the pilot to deal with emergencies continues to draw attention in aviation, with Airbus favoring the former while Boeing has favored the latter, up until MCAS.[6]

The automotive industry needs to consider the path to L5 very carefully. It might be worthwhile changing course, skipping L3 altogether (Ford, Waymo, and Audi had announced they would skip L3, but Ford has since reversed its position[8]), and designing a fully autonomous vehicle, bypassing the fraught nature of semi-autonomous controls.

Table 1. U.S. motor vehicle accidents for 2018.[9]

Year	Deaths	Vehicle Miles Traveled (VMT) (Billions)	Fatalities/ 100 Million VMT	Population	Fatalities/ 100K People
2018	36,560	3,174	1.13	327,200,000	11.17

Table 2. Projected death rates for autonomous vehicles.

Case	Safety relative to current manual benchmark	Failure Rate (per hour)	Annual Deaths caused by Control System (US)	Deaths/Day (US)
1	Same as	5×10^{-7}	36,560	100
2	10X better	5×10^{-8}	3,656	10
3	100X better	5×10^{-9}	365	1

Conclusion

The interest from both the industry and the driving public in autonomous vehicles is considerable and justified. But fielding technology that promises huge societal impact without a serious consideration of its dependability requirements is unsound. "Better than the average driver" is a particularly weak requirement. Engineers have proven they can do much better than that in other fields. Society needs to provide the incentive for them to do what needs to be done in the automotive domain. C

References

1. Automotive Information Sharing and Analysis Center (Auto-ISAC). *Best Practices Executive Summary.* (July 2016); https://www.automotiveisac.com/best-practices/
2. DeNisco-Rayome, A. Dossier: The leaders in self-driving cars. *ZDNet* (Feb. 1, 2018); https://www.zdnet.com/article/dossier-the-leaders-in-self-driving-cars/
3. Eliot, L. Discussing safety first for automated driving with Aptiv's Karl Iagnemma. *Forbes.* (July 19, 2019); https://www.forbes.com/sites/lanceeliot/2019/07/19/discussing-safety-first-for-automated-driving-with-aptivs-karl-iagnemma/
4. ISO/SAE 21434. Road vehicles—Cybersecurity Engineering; https://www.iso.org/standard/70918.html and https://www.sans.org/cyber-security-summit/archives/file/summit-archive-1525889601.pdf
5. Lala, J. History and future directions of mission- and safety-critical digital avionics. In *Proceedings of the AIAA Guidance, Navigation & Control Conference* (Aug. 2013), Boston, MA; https://doi.org/10.2514/6.2013-5206
6. Langewiesche, W. What really brought down the Boeing 737 Max? *New York Times Magazine.* (Sept. 18, 2019); https://www.nytimes.com/2019/09/18/magazine/boeing-737-max-crashes.html
7. Levin, A. and Beene, R. Max disasters fuel outcry over how FAA let Boeing self-certify. *Bloomberg Markets* (Dec. 3, 2019); https://www.bloomberg.com/news/articles/2019-12-03/max-disasters-fuel-outcry-over-how-faa-let-boeing-self-certify
8. Martinez, M. Ford rethinks Level 3 autonomy. Automotive News Europe. (Jan. 20, 2019); https://europe.autonews.com/automakers/ford-rethinks-level-3-autonomy
9. National Highway Traffic Safety Administration (NHTSA) Traffic Safety Facts Research Note. (Oct. 2019); https://www.nhtsa.gov/traffic-deaths-2018
10. SAE Standards News: J3016 automated-driving update. (Jan. 2019); https://www.sae.org/news/2019/01/sae-updates-j3016-automated-driving-graphic
11. *Safety First for Automated Driving.* Technical Report. (July 2019); https://www.aptiv.com/docs/default-source/white-papers/safety-first-for-automated-driving-aptiv-white-paper.pdf
12. USDOT Automated Vehicle Activities. (Apr. 2020); https://www.transportation.gov/AV
13. Wakabayashi, D. Self-driving Uber car kills pedestrian in Arizona, where robots roam. *New York Times* (Mar. 19, 2018); https://www.nytimes.com/2018/03/19/technology/uber-driverless-fatality.html

Jaynarayan H. Lala (jay.lala@rtx.com) is a Senior Principal Engineering Fellow at Raytheon Technologies, San Diego, CA, USA.

Carl E. Landwehr (carl.landwehr@gmail.com) is a Research Scientist at George Washington University and a Visiting Professor at University of Michigan, Ann Arbor, MI, USA.

John F. Meyer (jfm@umich.edu) is a Professor Emeritus of Computer Science and Engineering at University of Michigan, Ann Arbor, MI, USA.

This Viewpoint is derived from material produced as part of the Intelligent Vehicle Dependability and Security (IVDS) project of IFIP Working Group 10.4.