SAFECOMP 2018 PREPRINT

Practical Experience Report: Automotive Safety Practices vs. Accepted Principles

Philip Koopman

Carnegie Mellon University, Pittsburgh PA 15217, USA
koopman@cmu.edu

Abstract. This paper documents the state of automotive computer-based system safety practices based on experiences with unintended acceleration litigation spanning multiple vehicle makers. There is a wide gulf between some observed automotive practices and established principles for safety critical system engineering. While some companies strive to do better, at least some car makers in the 2002-2010 era took a test-centric approach to safety that discounted non-reproducible and "unrealistic" faults, instead blaming driver error for mishaps. Regulators still follow policies from the pre-software safety assurance era. Eight general areas of contrast between accepted safety principles and observed automotive safety practices are identified. While the advent of ISO 26262 promises some progress, deployment of highly autonomous vehicles in a non-regulatory environment threatens to undermine safety engineering rigor.

Keywords: Software safety, automotive, unintended acceleration

1 Introduction

Innocent people have died, been severely injured, or gone to jail because of defects or potential defects in computer-based automotive systems. With the deployment of self-driving cars, it is more important than ever to understand the gaps between theory and practice in automotive computer-based system safety.

This paper is based on the author's personal experiences with unintended acceleration (UA) litigation against car makers (Original Equipment Manufacturers, or OEMs) for 2000-2010 model year vehicles, and additional experiences with multiple recent military and commercial self-driving car (Autonomous Vehicle, or AV) safety assurance projects. These experiences include access to extensive sets of engineering documents, analysis of Electronic Throttle Control (ETC) source code, and vehicle testing to confirm identified safety vulnerabilities. These experiences have revealed common threads that encompass technical, business, regulatory, and litigation aspects of safety. While regulatory environments vary in other countries, the significant role that the United States (US) car industry and US legal system play in the automotive domain ensure that these factors will influence many cars produced worldwide.

Unlike other domains, conformance to international computer-based system safety standards is voluntary for US-sold vehicle OEMs and suppliers. Moreover, some

OEMs have not followed industry-specific guidelines such as the MISRA Software Guidelines [1], including vehicles that are the subject of two class action lawsuits. [2] at 30:21-25 and [3] at 78:15-79:15. (Note that [2] is a transcript from a death and injury case involving a vehicle of a type included in the corresponding class action lawsuit.)

The US permits OEMs to deploy vehicles that are self-certified to meet provisions of the US Federal Motor Vehicle Safety Standards (FMVSS). FMVSS regulations take the form of a test procedure approach originally intended to ensure that the normal safety-relevant functionality of pre-computer vehicles, such as braking capability, was adequate. While some simplistic failure modes such as detecting the complete loss of a functional subsystem are included, the test procedures are not intended to achieve any defined amount of software testing coverage, are not designed to detect non-deterministic faults, and do not demonstrate fault recovery from non-trivial computational faults. While vehicles commonly use some basic fault tolerance patterns such redundant CPUs for life-critical functions, it can be the case that redundancy and other fault tolerant computing techniques not used in accordance with accepted practices, such as dual-CPU designs with a single point of failure. [4]

At least one death has been officially declared to be due to automotive computer system malfunction [4], and there have been approximately 500 settlements for death and injury alleged to also be due to defective vehicle designs by the same OEM. [5] Another class action against a second OEM alleges similar issues. [3] Additionally, there are instances in which individuals have faced civil or criminal penalties for mishaps they claim were caused by vehicle malfunctions (e.g., [6]). Now that computers have life critical control authority, they must be considered as a credible potential cause of severe mishaps.

Electrified vehicles present additional risks because regenerative braking tends to disable the direct hydraulic connection between the brake pedal and friction-based brake pads. [7] (If this weren't the case, energy could be lost due to friction instead of being used to recharge the battery.) Some drivers have reported loss of brake effectiveness with these vehicles (e.g., [8]) which could potentially be caused by a software defect. Some litigation has involved reported symptoms consistent with such a defect. Increasing levels of autonomy raise the stakes further.

Table 1. Contrasting areas of safety principles and observed automotive practices.

Accepted Safety Principle	**Observed Automotive Safety Practice**
Evidence required to show safety	Evidence required to show defect
Safety argument	System-level functional test
Arbitrary failures	"Realistic" failures
Random failures expected	Non-reproducible failures are discounted
Blaming humans is a last resort	Driver error presumed
Engineering rigor and integrity level	All unsafe defects identified and fixed
Independent assessment	Self-certification
ALARP, etc.	Cost effective regulation

Table 1 identifies areas in which some observed OEM practices do not necessarily correspond with accepted safety principles. The scope of this table deals with vehicles produced with ETC in the 2002-2010 era from some Asia, US and European OEMs selling into the US market. It should be emphasized that some OEMs claim to follow accepted safety practices. And to be clear, the listed OEM practices should not be considered industry-accepted practices for making safe vehicles, but rather should be seen as areas in which some OEMs' observed practices fell short of meeting accepted safety practices. Based on personal experience in a variety of venues, it is clear that portions of the OEM and supplier ecosystem were still stuck in the pre-software safety engineering era at least up until the creation of ISO 26262 [9], and that adoption of that new standard is taking time.

2 Safety Principles vs. Automotive Safety Practices

2.1 Safety Arguments Aren't Specifically Required by Regulators

A general safety principle is that a system is not presumed to be safe until a mishap occurs, but rather must be demonstrated to be safe before deployment. Approaches to demonstrating safety are typically based on some sort of safety argument. That argument might be explicit (e.g., a GSN argumentation structure [10]), implicit in the form of having followed a suitable set of safety practices (e.g., [1]), or some mixture of the two. Common codified safety practices include the generic notions of a Safety Integrity Level (SIL), Design Assurance Level (DAL), or other risk-based approach to identifying and requiring a defined level of engineering rigor.

The US legal system, on the other hand, tends to emphasize the identification of defects. OEMs can attempt to defend themselves simply by asserting that their vehicle is safe because no bugs have been identified that lead to UA. [11] at 47:3-10. Injured parties and their experts typically must search for relevant bugs or other design defects such as single points of failure to support a vehicle defect argument.

US regulations do not require vehicles to have a safety argument beyond FMVSS compliance, although using one is not precluded. However, lack of following accepted engineering practices can be a contributing factor to legal outcomes, especially when considering negligence. Additionally, a pattern of mishaps can lead to a mandatory vehicle recall in some cases.

Some European vehicles in the 2000s adopted the E-Gas approach for electronic throttle control. ([12] is a newer, publicly available description.) In general, the approach involves a primary functional unit that performs control, and monitoring/checking units that disable engine power if a fault is detected. The suitability of this approach for life-critical applications depends upon adequate isolation between doer/checker levels and appropriate fault coverage. In some cases, independent UA mitigation is required, such as a vacuum pump to boost braking force independent of throttle position. The specification also describes required fault handling functionality.

2.2 Argumentation vs. Testing

While general safety principles require some sort of argument based in part on engineering analysis and rigor, the US regulatory system and much common practice is heavily based on vehicle-level testing. It is common for OEMs to practice non-software-specific techniques for fault analysis such as DFMEAs. [13] However, use of more advanced computer-based system safety techniques is uneven.

As previously discussed, the centerpiece of US automotive safety regulation is the suite of Federal Motor Vehicle Safety Standards (FMVSS). While some testing contemplates simplistic component fault models, FMVSS criteria generally do not involve design processes, code quality, or other accepted computer-based system safety considerations. For example, FMVSS 138 [14] fault injection covers a silent malfunction due to loss of component power in a tire pressure monitoring system. Similarly, US National Highway Traffic Safety Administration (NHTSA) investigations involve vehicle level testing and discussions with the OEM, but emphasize driver error as a cause of UA. For example, [15] blames the driver rather than the ETC for data samples showing a doubling of engine RPM and vehicle speed with unchanged accelerator pedal input.

2.3 Arbitrary vs. "Realistic" Faults and Failures

For safety critical systems, even a single bit flip or other small fault has the potential to cause a catastrophic mishap if not sufficiently mitigated. Well defined and expansive fault models such as transient faults and single event upsets are well known in the areas of safety and fault tolerant computing research. Arbitrary failures of computer-based system components must be considered when designing life-critical systems. [16] Moreover, there is an increasing body of confirmed reports of Byzantine (e.g., two-faced) faults occurring in real systems. [17] However, some OEMs do not embrace these accepted fault and failure models.

Automotive OEM safety analysis is often concerned with simplistic fault models such as electrical wires shorted to power supply voltages, open circuits, or computer crashes. Faults that are subjectively judged not to be "realistic" by designers are often dismissed. However, research has documented subtle real world faults and failures that defy designer intuition about fault realism. [18]

Any redundancy often relies upon self-diagnosis and simplistic fault detection mechanisms such as watchdog timers, heartbeats, and input port sanity checks. [4] Such simplistic redundancy management approaches offer only partial fault coverage, and permit dangerous fail-active behaviors. [19]

2.4 Failure Reproducibility

Transient faults and resulting failures are generally not reproducible upon demand in ordinary system operation, because the underlying causes can be comparatively infrequent, randomly occurring events. Fault injection experiments reveal vulnerabilities, but are routinely criticized in litigation for involving minor instrumentation modifica-

tions to vehicle software such as inclusion of a subroutine to flip memory bits upon command. Such modifications are then claimed to render fault injection results invalid due to involving a variation from the exact software image that would be in a production vehicle, or otherwise not being "realistic." [11] at 84:14-24.

Diagnostic gaps and undiagnosed failures are common. In some – but not all – cases, Trouble Not Identified (TNI) incidents can eventually be traced to systematic causes with sufficient detective work. [20] Despite less than complete diagnostic coverage, and substantial TNI rates, ETC malfunction is often inappropriately ruled out by OEMs or investigators when no Diagnostic Trouble Code (DTC) has been recorded. This is especially true when problems cannot be reproduced with the subject vehicle – even when a report is made by a source that many would consider credible, such as a dealership employee or police officer. [3] at 86:10-87:24.

Automotive safety struggles with non-reproducible faults. NHTSA tends to close investigations of non-reproducible faults rather than investigating potential software defects as root causes of mishaps. Similarly, OEMs can emphasize reproducible faults and undeniable trends of field data, rather than perceived "one-off" events, in part to avoid putting "the company out of business." [21]

2.5 The Driver Error Narrative

It is well known that humans are imperfect. It follows that the heart and soul of a typical UA legal defense is a claim of driver error, typically in the form of pressing the accelerator pedal instead of the brake pedal. Many publications, including those from NHTSA, repeat the refrain of driver error causing UA events. [22] However, these reports fail to consider computer system defects. Rather, reports conclude that in the absence of mechanical defects or concrete physical evidence of a vehicle malfunction the cause of a mishap must be drive error. Situations that provide truly compelling evidence to rule out drive error tend to be attributed to "unknown" causes.

While OEMs and NHTSA typically cite various reports in support of the pedal misapplication narrative, what data can be found on that specific failure mode tends to tell a different story. A pre-ETC analysis of 997 "reasons/excuses" for crashes found only one instance of "hit gas pedal instead of brake" – but 29 instances of "vehicle failure." [23] pp. 293, 296. Thus, contrary to the typical human error narrative, available data provides support for a finding that vehicles malfunction more often than humans press the wrong pedal.

Revisiting the Audi 5000 investigation report reveals that even the veritable poster child of human error producing UA provides incomplete support for the pedal misapplication narrative. Audi vehicle malfunctions produced up to 0.3g of un-commanded acceleration, having nothing to do with driver error. However, when such a UA event startled the driver, sometimes the driver would press the wrong pedal, resulting in a collision before there was time to self-correct in a tight-quarters situation. [24]

Pedal misapplication issues are complicated by problems with data recording strategies, such as potentially missing driver actions due to under-sampling. [15] Moreover, data recordings can be untrustworthy to the extent they rely upon suspect data being provided by the same computer that is potentially causing the UA.

2.6 Engineering Rigor

Developing naked, undocumented code with no substantive safety process can reasonably be expected to result in defects that could cause a catastrophic loss event for life critical systems. This can create a fear that developers will be criticized for the smallest of imperfections. However, the remedy for this fear is well understood: use an accepted safety approach. If nothing else, a successful independent assessment provides an argument in defense of allegations of negligence. However, a negative assessor report can appear to be adverse in litigation. [3] at 78:15-78:15.

Some automotive designers adopted model-based design during the 2000-2010 timeframe. This type of approach can provide tool support for certified code generation and formal proofs of correctness for some aspects of system operation. However, more than this is required for safety, and use of this type of tooling does not by itself ensure good design quality. The two class action cases discussed in this paper did not make any apparent use model based design for the code in question.

2.7 Certification and Deployment of Autonomous Vehicles

Independent assessment of safety standard conformance has been possible for many years in the automotive industry. However, current automotive regulations only require assessment against FMVSS test regimes. The future of AVs currently promises more of the same. A first draft AV policy [25] encouraged some level of accountability for safety arguments via a self-certification signature sheet. However, a later version takes a "non-regulatory" approach to safety, making even self-certification entirely optional for AVs. [26] Current US federal regulatory efforts emphasize modifications or waivers of FMVSS test regimes to accommodate AVs.

Of significant concern in AV deployment is the usual argument for doing so: human drivers make avoidable mistakes; computers won't make those mistakes; therefore computers will be safer drivers than humans. There is insufficient field data and no robust technical public safety argument upon which to base an assertion that AVs have even achieved safety parity with an "average" human driver (whatever that might actually mean, noting that impaired drivers are part of the human driver population). Perhaps AVs will simply make *different* mistakes. Ensuring AV safety is complicated by the use of novel technologies such as machine learning. [27]

Two vendors have commendably published safety brochures. [28][29] No vendors currently claim rigorous, independently assessed safety arguments.

3 Regulatory and Litigation Considerations

3.1 Cost Effectiveness of Safety Assessment

Accepted safety practices require reducing risk to an acceptably low level, e.g., As Low as Reasonably Practicable (ALARP). However, US government agencies are required to justify that all new regulations, including safety regulations, are cost effective. The existing pedal misapplication narrative surrounding UA makes it difficult to

introduce new software safety regulations to avoid software defects, because such defects have not been officially blamed for many mishaps. If there is no apparent carnage from unsafe software, it is difficult to cost-justify improving software safety. However, new laws can create stronger safety requirements without cost justification.

The litigation aspect of cost effectiveness is a bit different. Generally, the questions asked are whether accepted engineering practices were followed, and whether a reasonable alternative design approach would have prevented a mishap from occurring. However, a defect must first be identified before those questions are asked, and generally some sort of loss or legal violation must occur before legal action can be taken.

3.2 Source Code Availability

Source code is generally unavailable for inspection unless a very large litigation effort is mounted. Government regulators do not have access to source code, nor do any outside assessors unless the OEM decides to voluntarily grant access. Even if litigation source code access is granted, it is often done under onerous conditions such as via a dedicated non-networked secure room with a metal detector wanding procedure before entrance. In one case, a judge found that OEM "misrepresentations caused Plaintiffs to incur unnecessary costs" due to requiring overly burdensome source code security measures. [30] All things considered, source code analysis can easily turn into a million-dollar-plus effort including the cost of litigating to gain access, the cost of operating a secure room, and expert witness costs. This makes source code analysis impractical for most litigation, especially criminal defense, unless it can piggy-back on a class action lawsuit that has deep pockets financial backing.

The expense and difficulty of source code analysis provides a perverse incentive for poor code quality, skimpy design information, and opaque configuration management practices. The more difficult to understand the software system is, the more difficult and expensive it will be for experts to access it and identify specific defects that could have caused UA or other dangerous vehicle behaviors.

3.3 The Importance of Academic Rigor in Publication

Academics need to be aware that litigation uses peer-reviewed academic papers as evidence to support expert testimony. Even a well-intentioned paper that reaches a flawed or poorly stated conclusion can do significant damage to practical safety if a lawyer can find a way to interpret it as providing protective cover for an unsafe system. Researchers and reviewers should be mindful of ways in which a paper might be used to support an opinion that accepted safety practices are deficient unless that is truly the finding of the research data. A particularly important point is that old techniques should not be identified as defective simply because new techniques are better. Studies should disclose threats to validity so that conclusions are not applied in inappropriate situations. Finally, reviewers and editors should ensure that authors who attempt to discredit previous publications fully disclose potential conflicts of interest that might potentially result in bias, such as involvement in pending litigation adverse to the previous publication's findings or authors. [31]

4 Conclusions

Automotive-specific safety guidelines and standards have existed for more than two decades. Yet adoption is not required, and not is universal. Recent findings of industry cover-ups regarding sticky gas pedals, floor mats, ignition switches, air bags, and emission defeat devices do not inspire confidence. One can hope that the significant costs paid by OEMs for these transgressions will motivate better behavior in the future. Litigation historical outcomes notwithstanding, it remains to be seen whether AV designers will adopt robust safety engineering practices, or will succumb to pressure and take shortcuts in the rush to market.

While it would be best if all OEMs actually adopted well understood accepted safety practices, a more pragmatic approach is to perform research that will meet the automotive industry where it is instead of where it should be. To that end, additional work on the following topics could help improve practical automotive safety (this list should not be interpreted as criticism of currently accepted safety practices):

- Studies that explicitly differentiate between driver error and computer faults
- Studies that measure how well specific safety techniques reduce mishap risk
- Fault injection techniques tailored to production vehicle deployment
- System-level testing approaches that validate safety
- Safety measurement approaches suitable for FMVSS test procedure codification
- Forensically valid automotive data recorders
- AV-specific safety validation (e.g., machine learning safety validation)
- Better understanding of the factors that support a robust safety culture

More generally, anything that the safety community can do help educate regulators, lawmakers, and non-specialist automotive practitioners appreciate the importance of adopting safety techniques proven in other domains can also help.

Threats to validity: Reported experiences are based on previous-generation vehicle designs due to the retrospective nature of the litigation and regulatory system. There is a significant variation in OEM attitudes and practice of safety, and certainly some OEMs try hard to adopt and even go beyond basic accepted safety practices.

Disclosure: The author is involved in ongoing litigation concerning multiple OEMs, including Toyota and Ford, and is a principle in an autonomous vehicle safety company. He is not a lawyer. No external support funded this research.

References

1. MISRA, Development Guidelines for Vehicle Based Software, Nov. 1994.
2. Bookout v. Toyota Trial Transcript, 11 Oct 2013, PM. https://goo.gl/MP8w3w
3. Charles Johnson et al. v. Ford Motor Company, US Dist. S. WV, Huntington, 3:13-CV-06529, 1 Feb 2018 PM. (Lawyer summaries of expert testimony and evidence.)
4. Koopman, P., A Case Study of Toyota Unintended Acceleration and Software Safety, Carnegie Mellon University, 18 Sept. 2014, presentation slides.

5. Kennedy, J., “Toyota has reached deals in 496 cases in acceleration MDL,” Law360, 15 Nov. 2017. https://goo.gl/T4TaLs
6. Manganis, J., “Cop’s fatal-crash trial underway; defense appears to abandon long-touted ‘sudden acceleration’ theory,” Salem News, 17 Mar. 2008. https://goo.gl/jiZ9rN
7. Toyota, 2005 Prius Repair Manual (RM1130U), page 05-951.
8. Marosi, R. & Olivarez-Giles, N., “Runaway Prius driver: I was laying on the brakes but it wasn’t slowing down.” 10 Mar. 2010. https://goo.gl/aZK7BM
9. ISO: Road vehicles-Functional Safety-Management of functional safety, ISO 26262, 2011.
10. GSN Community Standard Version 1, Nov. 2011.
11. Bookout v. Toyota Trial Transcript, 22 Oct 2013, PM. https://goo.gl/hh47vg
12. EGAS Working Group, Standardized E-Gas Monitoring Concept for Gasoline and Diesel Engine Control Units, Version 5.5, 2013.
13. SAE, Potential Failure Mode and Effects Analysis in Design (Design FMEA), J1739_200901, 15 Jan. 2009.
14. GPO, Section 571.138, Standard No. 138; Tire pressure monitoring systems. 49 CFR Ch. V (10-1-11 Edition).
15. NHTSA, Denial of a petition for a defect investigation, Federal Register Vol. 80, No. 93, 14 May 2015, pp. 27835-27844.
16. Lala, J., Harper, R., "Architectural principles for safety-critical real-time applications," Proceedings of the IEEE, Volume 82, Issue 1, Jan. 1994, pp. 25-40
17. Driscoll, K., Hall, B., Sivencrona, H. & Zumsteg, P., “Byzantine fault tolerance, from theory to reality,” SAFECOMP 2003, LNCS 2788, pp. 235-248.
18. Driscoll, K., Real System Failures, 2012. c3.nasa.gov/dashlink/resources/624/
19. Hammett, Design by extrapolation: an evaluation of fault-tolerant avionics, 20th Conference on Digital Avionics Systems, IEEE, 2001.
20. Thomas, D. et al., The ‘trouble not identified’ phenomenon in automotive electronics, Microelectronics Reliability, pp. 641-651, 2002.
21. Gladwell, M., “The engineer’s lament: two ways of thinking about automotive safety,” The New Yorker, May 4, 2015.
22. Lococo, K. et al., Pedal Application Errors, DOT HS 811 597, Mar. 2012.
23. Wierwille, W., et al., Identification and evaluation of driver errors; overview and recommendations. Federal Highway Administration; McLean, VA, FHWARD-02-003, 2002.
24. Walter, R., et al., Study of mechanical and driver-related systems of the Audi 5000 capable of producing uncontrolled sudden acceleration incidents, DOT-TSC-NHTSA-88-4, Dec. 1988.
25. US DoT, Federal Automated Vehicles Policy: Accelerating the next revolution in roadway safety, Sept. 2016.
26. US DoT, Automated Driving Systems 2.0: A Vision for Safety, Sept. 2017.
27. Koopman & Wagner, Autonomous Vehicle Safety: an interdisciplinary challenge, IEEE Intelligent Transportation Systems Magazine, Spring 2017, p. 90-96.
28. Waymo, On the Road to Fully Self-Driving, 2018. https://goo.gl/3GwP2T
29. GM, 2018 Self-Driving Safety Report. https://goo.gl/2d5PTM
30. Charles Johnson et al. v. Ford Motor Company, US Dist. S. WV, Huntington, 3:13-CV-06529, order granting sanctions, 12/27/2017.
31. Koopman, P., Letter to Editor, IEEE Consumer Electronics Magazine, Jan. 2018.