How Safe is Safe Enough? Acceptable Safety Criteria From an Engineering and Legal Perspective

by Martin S. Chizek
Orlando, Florida

Manufacturers have a vested interest in the safety of their customers, and in protecting their reputation for producing safe products. An additional incentive to produce safe products is avoiding liability when their product is involved in an accident or mishap that results in personal injury and/or property damage. While it is often said that one must never compromise on safety, the fact remains that any product must necessarily be a balance between the level of safety desired and the cost and performance impact of achieving that level of safety. The product manufacturer must make a determination: Is this product (or technology) acceptably safe within the context of current consumer expectations as well as the legal/regulatory framework? Is the residual risk tolerable? This paper presents a methodology to address those questions by reviewing the publicly available information of a recent automotive product liability case, and evaluating whether the product design met current legal and safety engineering best practices.

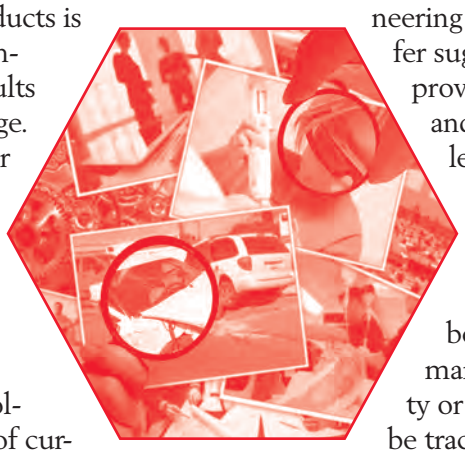

Introduction

It rarely makes sense to attempt to make a product or system *absolutely* safe and risk-free. On an economic level, the manufacturer should strive to achieve a level of risk that equates the incremental benefits of greater safety with the incremental costs. From an engineering perspective, this generally means implementing a design safety program that identifies product hazards and reduces them to an acceptable level through design best practices and user warnings within the constraints of program cost and schedule. If the product is later involved in personal injury or property damage, it will be adjudged *ex post* by the legal system as to whether the product contained a defect that made it unreasonably dangerous and was the legal cause of the injury or damage.

This paper will also explore the question of "how safe is safe enough" from both a legal and engineering viewpoint, particularly in the automotive industry. First, we list the elements of a system safety process common to most industries. Then, we explore the legal rules and court decisions on how a product design is determined to be "reasonably safe." Next, we review the jury results of a recent automotive "unintended acceleration" case, and juxtapose the jury findings against the identified engineering and legal best practices. Finally, we offer suggestions to auto manufacturers for improvement in the areas of safety standards and liability risk reduction in the current legal and regulatory environment.

Achieving "Safe Enough" from an Engineering Perspective

System safety engineering principles can be applied to any activity to reduce or manage the risk of harm to people, property or the environment. These principles can be traced back to early military and aviation standards [Ref. 1]. Most industries use a safety engineering process comprised of common elements, such as the following:

Determine Safety Goals and Requirements: Establish general safety objectives, specific safety performance requirements, and risk levels considered acceptable for the system. Risk levels can be defined in terms of a mishap risk category, an overall system mishap rate, demonstration of controls required to preclude unacceptable conditions, or satisfaction of specified standards and regulatory requirements. Quantitative requirements may be expressed in terms of either risk, or the probability or frequency of a given mishap severity category.

Identify Applicable Regulations and Standards: Determine if the product is regulated by federal or state regulations or agencies. If the company is going to market the product in the European Union (E.U.), review applicable E.U. Directives. Next, determine the voluntary consensus standards most applicable to the product. If a U.S. standard is not available, there is almost certainly an International Standards Organization (ISO) standard available.

Identify and Track Hazards: Identify hazards through a systematic analysis process that includes system hardware and software, system interfaces (including human interfaces), and the intended and foreseeable environments of use. Environments to consider should include intended (normal) operation, malfunctions and reasonably foreseeable misuses of the product.

Analyze Hazard(s) and Assess Risk: Assess mishap severities and probabilities or frequencies for each hazard across each mode of operation. Use accepted hazard analysis techniques to identify early in the life cycle those risks that can be eliminated by design and those that must undergo mitigation by other controls to reduce risk to an acceptable level. Software and programmable logic should be subjected to current best practices for safety-critical software [Refs. 1 & 2].

Reduce Risk: Prioritize hazards so that controls and risk mitigations focus on the most serious hazards according to the mishap risk potential they present. Implement mitigations according to the order of precedence, consisting of eliminating the hazard through design selection or inherently safe design measures, incorporating detection and safety devices to guard against the hazard, and/or providing applicable warnings and instructions on how to avoid potential hazards while using the product.

Verify and Validate Risk Reduction: Verify the implementation and validate the effectiveness of all selected risk mitigation measures through appropriate analysis, testing, demonstration or inspection. Test or demonstrate safety-critical components and functions to establish the design's margin of safety.

Determine Residual Risk Acceptance: All reasonably foreseeable hazards must be identified, evaluated and mitigated to a level compliant with applicable laws, regulations and company policy before the product is released for customer use. If available, a risk acceptance authority should determine whether the mishap risks have been reduced to an acceptable or tolerable level, and either accept the residual mishap risk or require further risk reduction.

Achieving "Safe Enough" from a Legal Perspective

Product liability law is intended to encourage manufacturers to produce safe products by subjecting them to liability when their product falls below an acceptable level of safety; i.e., when the product is *not* safe enough. A manufacturer owes *a duty of care* to consumers to make the product reasonably safe for use, and to provide warnings about any dangers inherent in the product.

While the law does not require a product to be completely risk free, it does require that a company in the business of producing specific types of products exercise a *standard of care* that is reasonable for those who are experts in designing and manufacturing such products. As explained in a 2007 federal court case, "While a manufacturer has a duty to design a product that is *reasonably safe for its foreseeable use*, it is not required to design the best possible product. Proof that technology existed, which if implemented could feasibly have avoided a dangerous condition, does not alone establish a defect" [Ref. 3]. The *state of the art* at the time of manufacture will usually be considered when assessing whether a reasonable standard of care was achieved. This includes current industry customs and standards, and the technological feasibility of safety devices and measures. Proof that a product could not have been made safer, under the practical technological feasibility existing at the time of manufacture, is conclusive evidence of due care [Ref. 7].

Negligence: A liability suit based on the *negligence* theory generally alleges that the seller or manufacturer (the defendant) breached a duty of care to the injured party (the plaintiff) by failing to eliminate a *reasonably foreseeable risk of harm* associated with the product that caused the harm. Such suits typically claim negligently defective design, negligent manufacture of the goods (including inspection and testing) and/or negligent failure to provide adequate warnings of hazards or defects.

Strict Liability: Most states have adopted some form of strict product liability that does not require showing negligence. The most common version of strict product liability sets out that "One who sells any product in a *defective condition unreasonably dangerous* to the user or consumer or to his property is subject to liability for physical harm thereby caused to the ultimate user or consumer, or to his property." This rule applies even though "the seller has exercised all possible care in the preparation and sale of his product" [Ref. 4]. Therefore, liability is not based on whether a product was safe or unsafe in any sort of absolute sense, but upon whether a product was *reasonably* safe. When determining whether a manufacturer should have made the product safer, most jurisdictions will require evidence that an *existing reasonable alternative design* would have prevented the harm that occurred and made it reasonably safe [Ref. 5].

Design Defect: This occurs where the design of the product makes it *unreasonably dangerous for its intended purpose*. Design defect cases frequently involve such factors as (i) the magnitude or severity of the foreseeable harm, (ii) industry practices at the time the product was manufactured, (iii) the state of the art of existing scientific and technical knowledge at that time and (iv) the product's compliance or noncompliance with government and industry safety regulations and standards. Most jurisdictions employ a *risk-utility analysis* when weighing these factors, including the design's social utility, and the effectiveness and cost of alternative safer designs. Other courts use the "consumer expectations test" that focuses on the reasonable expectations of the consumer or purchaser, and requires the court to determine whether the product "is more dangerous than an ordinary consumer would expect when used in an intended or reasonably foreseeable manner" [Ref. 6].

Compliance with Regulations and Standards: In most jurisdictions, violation of a mandatory government

regulatory standard results in an automatic finding of negligence. An exception is when a federal statute or regulation requires that a product be manufactured with a particular design or warning, or when the product has received close scrutiny and approval by a designated federal agency. In these cases, the mandated or federally approved design or warning is not considered defective or inadequate under state law via the concept of *federal preemption*. However, compliance with applicable laws and regulations is not, for most products, an absolute defense in a product liability case. The plaintiff may argue that the manufacturer should have exceeded laws and regulations pertaining to safety. Similarly, compliance with industry standards and certifications such as Underwriters Laboratories may provide evidence that the product is reasonably safe, but the plaintiff can argue that a manufacturer exercising a reasonable standard of care would have exceeded the standards.

Pre-Trial Discovery: Each party to a lawsuit is entitled to request that other parties produce documents that are in their possession or control. Any document arguably relevant to the case is subject to production. A plaintiff's lawyer prosecuting a significant products liability claim may, as a matter of course, ask for all documents [Ref. 7]:

- Related to the design of the product
- Related to specifications, and change in the specifications for the product
- Related to quality control procedures used in the product's manufacture
- Related to the source of the components of the product, or at least of those components involved in the accident
- Related to marketing, promotion and advertising of the product
- Related to the organization of the manufacturer, and detailing who was responsible for decisions about the product

Expert Witnesses: In many cases, the plaintiff in a complex product liability suit cannot prove that a product was defective without expert testimony. Theoretically, the sole function of an expert witness is to educate the finder of fact (jury) about pertinent matters beyond the competence of laymen. In recent years, reliance on expert testimony at trial has increased, and experts are allowed to go further toward giving an opinion on the ultimate issues of the case than once was allowed. If an expert is retained by a claimant's lawyer to examine a product and finds nothing wrong with the product, the lawyer may dismiss that expert and seek another. The existence of the unfavorable opinion of the first expert is protected by the work-product doctrine. Because neither party is under any compulsion to use an expert not entirely favorable to his case, experts often act as advocates rather than objectively educating a jury [Ref. 7].

Burden of Proof: In most product liability suits, the burden of proof that applies is called "a preponderance of the evidence," and requires the jury to return a judgment in favor of the plaintiff if it is shown that a particular fact or event was *more likely than not* to have occurred. This is usually interpreted as requiring a finding that at least 51 percent of the evidence supports the plaintiff's allegations.

Damages: In addition to compensatory damages, punitive damages may be awarded to the plaintiff if it is demonstrated that the defendant showed a *reckless or conscious disregard* for consumer safety. Knowing violations of safety standards, inadequate testing and manufacturing procedures, and failure to warn of known dangers may be evidence of such a disregard for safety [Ref. 7]. For punitive damages to be awarded, the burden of proof is generally elevated to a higher standard called "clear and convincing evidence" — or that a particular fact is substantially more likely than not to be true.

Toyota Unintended Acceleration Case Study

The National Highway Transportation Safety Administration (NHTSA) has investigated complaints of vehicles exhibiting unintended acceleration (UA) for decades, some of the most serious involving Toyota vehicles. NHTSA had concluded that these occurrences were the result of the driver accidentally pressing the accelerator pedal instead of the brake; floor mats and other obstructions that entrap the accelerator pedal; and damaged or malfunctioning mechanical components, such as broken throttles, frayed cables and sticking accelerator pedal assemblies. However, NHTSA continued to receive reports of UA where driver error and mechanical failure seemed unlikely.

In 2010, NHTSA enlisted the National Aeronautics and Space Administration (NASA) to investigate the potential for vulnerabilities in the 2005 Toyota Camry Electronic Throttle Control System (ETCS) because it had the highest rate of reported UA events. NASA identified the critical functions of the ETCS, examined how the electronics system was designed and implemented to guard against failures, and whether it responded safely when failures did occur. The system's design and implementation were specifically assessed for circumstances in which a UA failure could occur and go undetected so as to bypass system fail-safe responses. The following information is taken from the NASA Report unless otherwise noted [Ref. 8 & 9].

Toyota ETCS Description: As shown in Figure 1, the ETCS is composed of an accelerator pedal assembly, a throttle body assembly and an Engine Control Module (ECM). The ECM contains two Central Processor Units (CPUs), throttle motor control drive circuitry, a power supply and inputs from other functions. The prime sensors VPA1 and VTA1, and the Main CPU, control the intended throttle opening. The secondary sensors VPA2

and VTA2, and the Sub CPU, are used to validate consistent sensor data and a properly operating Main CPU. Both CPUs must agree that the throttle motor should be engaged in order for the throttle motor to drive the throttle valve open. A failure in a single pedal or throttle sensor will cause the associated CPU to declare a fault and transition to "limp home" mode. A failure in more than one sensor will cause the associated CPU to disable the throttle.

In Figure 2, the Main CPU uses a real-time operating system (RTOS) based on the OSEK standard for distributed control units in vehicles, which is supported by AUTOSAR[1]. The operating system is based on the execution of tasks, each with a fixed and statically assigned priority. Primary functions are analog and digital sensor input, control output and functional processing of the throttle valve, fuel injectors and ignition timing. Software modules include the following:

- The *Pedal Command Function* converts two accelerator pedal position sensor inputs into a desired throttle angle command. The pedal command is sensed by the software as the difference between the pedal released position and the pedal pressed position. The pedal function also contains "limp home" mode logic to limit the throttle in the event of pedal sensor failures.
- *Idle Speed Control* sets the throttle angle to achieve the desired idle speed, which is limited by software to a maximum of a 15-degree relative opening.
- *Throttle Control* has authority for throttle control functions, including Pedal Command and Idle Speed Control. The throttle command drives the throttle motor that rotates the throttle valve against its return springs. The throttle valve position is sensed by two sensors, which provide closed-loop feedback to the throttle motor driver. If a throttle valve position sensor fails, power is cut to the throttle motor.

The Sub CPU (also know as the Monitor CPU) performs hardware sensor input data and limit checks, self-diagnostic checks on the Main and Sub CPUs, and sets fault codes that can disable the throttle motor power. A heartbeat/watchdog exchange between the Main CPU and Sub CPU detects major CPU failures and can reset the CPUs, thereby disabling the throttle motor in hardware. Both CPUs use non-volatile ROM for software code and volatile Static RAM (SRAM). The SRAM is protected by a single error detect and correct, and a dou-

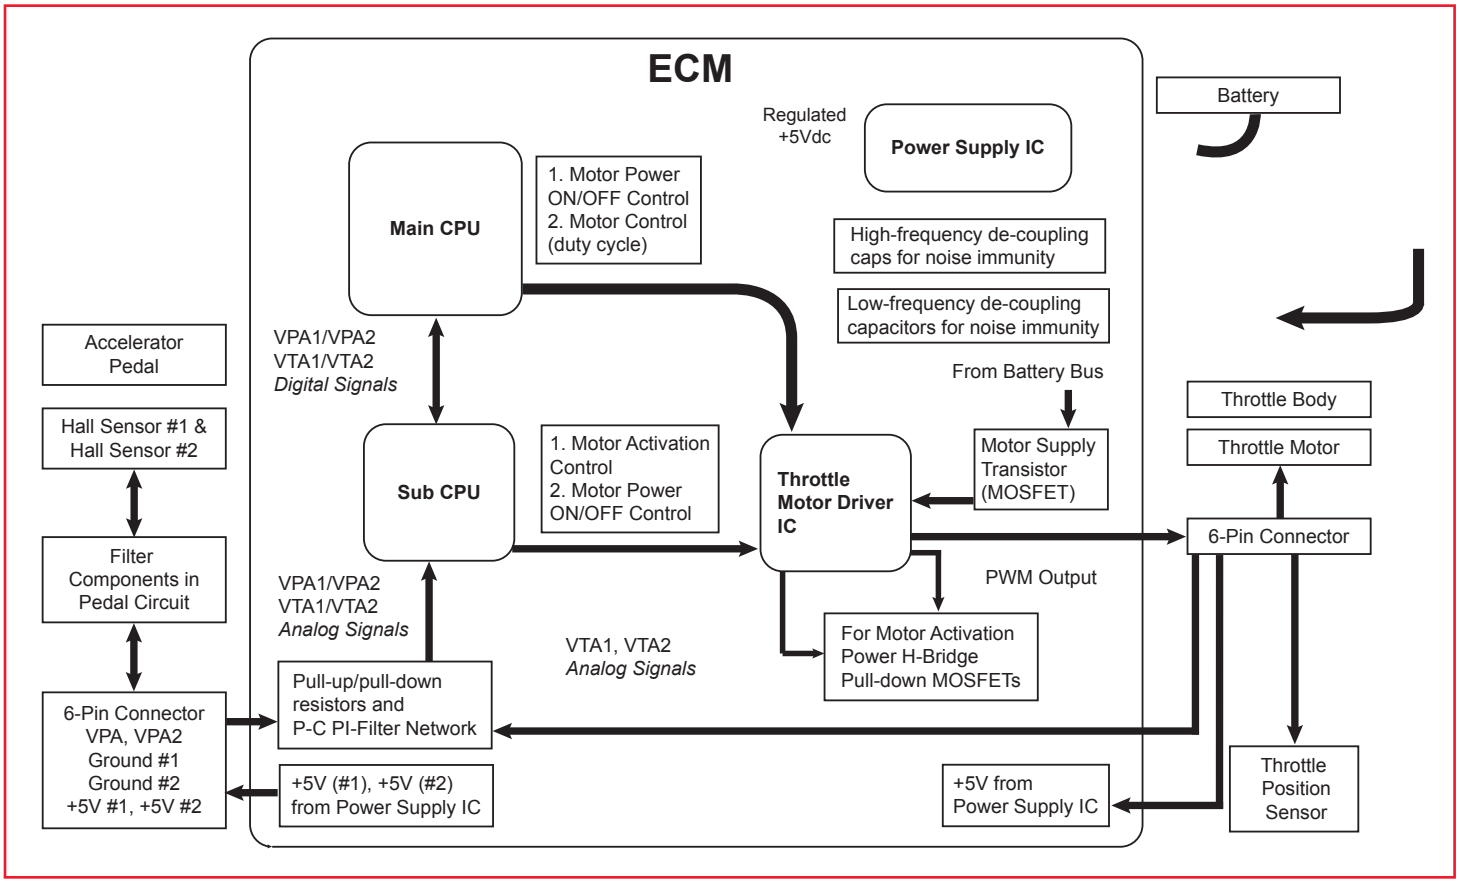

Figure 1 — ETCS Throttle Control Block Diagram [Ref. 16]

[1] OSEK is a standards body that has produced specifications for a standard software architecture and operating system for automotive embedded systems. AUTOSAR (AUTomotive Open System ARchitecture) is a partnership of automotive-related companies for establishing an open and standardized software architecture for automotive electronic control units.

ble error detect hardware function performed by error detection and correction (EDAC) logic.

The Main and the Sub CPU use hardware watchdog timers that are initiated at start-up and require constant re-initiation by software. If a watchdog timer expires or an abnormal condition occurs, the CPU hardware is reset and restarts. A subset of software data is protected by implementing software data mirroring. The data is written to two separate locations, which are then cross-checked when read. If the check fails, a default value is used. This data mirroring protects data from being overwritten if a stack or buffer overflow occurs. Processor and memory protection against Single Event Effects (SEE) includes EDAC on memory, data mirroring for critical variables, and watchdog timer and heartbeat functions between the two processors. The Electronic Fuel Injection (EFI) module employs fuel cut and ignition timing to mitigate the consequences of unintended throttle opening due to the failure of sensors, CPU or a mechanically stuck-open throttle valve.

The NASA Investigation: NASA identified two scenarios through Failure Modes Effects Analysis (FMEA) as having at least a theoretical potential to produce UA characteristic of a large throttle opening: (1) a systematic failure of software in the Main CPU that goes undetected by the Monitor CPU and (2) two faults in the pedal position-sensing system that mimic a valid acceleration command [Ref. 8]. NASA investigators used multiple tools to analyze software logic paths and to examine the programming code for paths that might lead to the first postulated scenario. While the team acknowledged that no practical amount of testing and analysis can guarantee that software will be free of faults, it reported that extensive analytic efforts uncovered no evidence of problems. To examine the second postulated scenario, the team tested numerous potential software and hardware fault modes by using bench-top simulators and by testing actual vehicles involved in reported cases of UA, including tests for electromagnetic interference. None of the testing could produce inadvertent acceleration indicative of a large throttle opening.

In 2011, NASA reported finding no evidence of Toyota's ETCS being a plausible cause of unintended acceleration events, and further concluded that the ETCS could not disable the brakes so as to cause loss of braking capacity, as often reported by drivers experiencing UA. Not having produced evidence of a safety-related defect in Toyota's ETCS, NHTSA elected to close its investigation into this system as a suspect cause of reported cases of high-power UA, and stood by its earlier conclusions attributing these events to pedal misapplication, floor mat entrapment and mechanical sticking. However, questions persisted as to whether Toyota's ETCS technology was to blame, particularly after media reports of more cases of Toyota vehicles exhibiting unintended acceleration, some involving fatalities.

"In 2013, plaintiffs sued Toyota on product liability and wrongful death theories claiming to have been injured in a 2007 car wreck involving their 2005 Camry. The plaintiffs claimed that the Camry accelerated unexpectedly when exiting a highway off-ramp and ran through an intersection and into an embankment. They claimed that the vehicle accelerated unexpectedly because of a defect in the car's electronic throttle-control system."

The Bookout v. Toyota Litigation: In 2013, plaintiffs sued Toyota on product liability and wrongful death theories claiming to have been injured in a 2007 car wreck involving their 2005 Camry. The plaintiffs claimed that the Camry accelerated unexpectedly when exiting a highway off-ramp and ran through an intersection and into an embankment. They claimed that the vehicle accelerated unexpectedly because of a defect in the car's electronic throttle-control system. The plaintiffs retained expert witnesses (Barr Group) to testify about Toyota's defective safety architecture and software defects [Ref. 10].

Expert Witness Testimony: Barr testified that they found what the NASA team sought, but couldn't find: "a systematic software malfunction in the Main CPU that opens the throttle without operator action, and continues to properly control fuel injection and ignition" that is not reliably detected by any fail-safe. The jury was told "it was more likely than not" that "Task X"[2] death caused the accident. The plaintiff's theory of the case was that undetected memory corruption in the Main CPU caused the unintended acceleration of the plaintiff's vehicle through a sequence of events:

[2] "Task X" apparently referred to a proprietary Toyota subprogram or thread that periodically executed on the Main CPU of the ETCS. Among its functions was to determine the correct throttle angle setting (how far open the throttle should be), based on how hard the driver is pressing on the accelerator pedal (as well as other factors). Task X determined the current accelerator pedal position and set a throttle angle variable accordingly, which was then used by another software task to set the throttle to the angle specified in that variable. Task X also contained the fail-safe modes, including "limp home" mode. [Refs. 10 & 11].

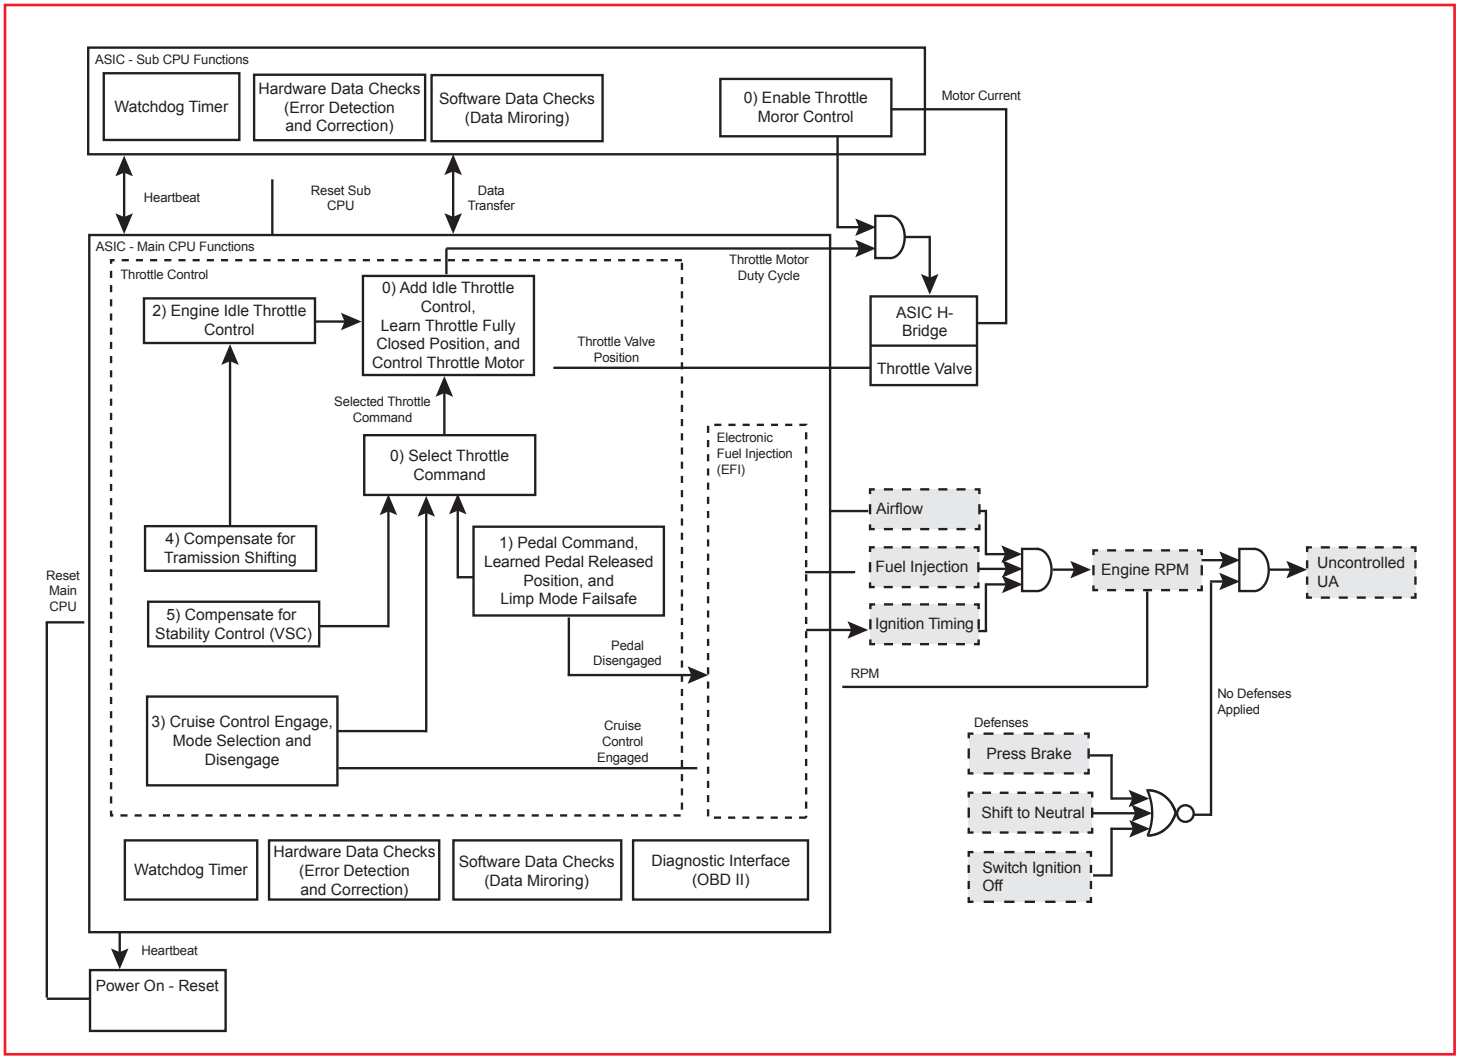

Figure 2 — ETCS Software Functions and System Safety [Ref. 8].

1. The bit corresponding to Task X in the operating system data structure "flipped" from "one" to "zero," resulting in the death of Task X.
2. At the time of this bit flip, the throttle angle variable maintained by Task X contained a large value, corresponding to an open throttle. Because Task X stopped executing, the throttle angle variable was stuck at the last computed throttle command value and the throttle remained open. (An alternative theory was that a second memory corruption caused the throttle angle variable to be overwritten with a large value.) When the plaintiff removed her foot from the accelerator pedal, there was no effect; the Throttle Control task continued to drive the throttle motor (and thereby the engine) at open throttle.
3. When the plaintiff stepped on the brake, the Brake Echo Check in the Monitor CPU did not correctly detect the death of Task X and force the throttle to idle. Because the throttle remained open, the driver was unable to stop the vehicle by braking. The Idle Mode Fuel Cut fail-safe did not work because it, too, was part of the Task X thread that had stopped executing.

Trial Results: In October 2013, the jury rejected Toyota's defense that the crash was caused by driver error, found Toyota liable for the 2007 highway crash and awarded $1.5 million each to the two plaintiffs. The jury also found that Toyota acted with reckless disregard for the rights of the plaintiffs, and was poised to award punitive damages in the second phase of the trial. However, the parties reached an undisclosed settlement before this occurred. Toyota subsequently agreed to pay $1.1 billion to settle with a class of roughly 23 million customers over recalls for defects in its vehicles that caused sudden unintended acceleration [Ref. 12].

Analysis of Alleged ETCS Defects

Each major assertion of the expert witness is summarized here, followed by observations on how the finding supports the legal elements and theories put forward by plaintiff's counsel. The plaintiff's theory of the case is also examined.

Claim of Defective Software Design: Toyota's ETCS source code is defective, of unreasonable quality and contains bugs that can cause unintended acceleration. Some critical variables are not protected from corruption by mirroring in a second location (including target throttle angle global variable). There is no hardware protection against bit flips. Stack overflow can occur, resulting in memory corruption. Toyota did not follow all coding guidelines, which could result in "bugs."

- Barr used the relevant legal terminology to allege a design defect (source code) that made the product (ETCS) unreasonably dangerous for its intended purpose, and established the causation element (resulting in unintended acceleration). Barr asserted a failure to follow industry practices at the time the product was manufactured (mirroring critical variables), failure to use the state of the art of existing scientific and technical knowledge at that time (protection from bit flips and memory corruption), and failure to comply with industry standards (MISRA C coding guidelines) [Ref. 13].
- **Claim of Defective Hardware Design:** Toyota's watchdog supervisor design is defective and unreasonable, and could not detect the death of Task X. The Monitor CPU did not detect all Main CPU malfunctions, and was not designed fail-safe for UA or task death. The Throttle Control and the fail-safe modes were all in Task X; i.e., in the same fault containment region. Reasonable design alternatives were well known. The Monitor CPU could have included a proper UA defense: *IF* (driver is braking and throttle is not closing) *THEN* reset ECM. Per-car cost to add this safety feature is $0.00 (it's just bits).
- Design defects were alleged (watchdog failed to detect Task X latch-up; a single point of failure, or SPF, existed by having Throttle Control and fail-safe modes in the same software thread). An existing reasonable alternative design for the Monitor CPU is suggested, which would have prevented the harm that occurred and made the ETCS reasonably safe. Barr alluded to Toyota's risk-utility analysis, asserting that a safer design was basically cost free.

Claim of Defective Software Process: Toyota had a defective software process that contained unreasonable single points of failure and failed to exercise a safe standard of care for software. FMEA was deficient or incomplete, since SPFs were present in a safety-critical system. Toyota didn't perform code reviews on the Main CPU and used a non-standard OSEK. There was no EDAC protection against hardware bit flips because it generally costs less to make memory chips without EDAC. The Monitor CPU software was a vendor-supplied item, and Toyota failed to gain access to the source code and review it.

- Software defect was alleged (software with SPF). Barr asserted that Toyota failed to comply with the standard of care of a reasonable manufacturer of safety-critical control systems. Barr again maintained that standards were not followed, and that an alternative safer design was available but not selected due to a risk-utility analysis that favored low cost over safety.

Plaintiff's Theory of the Case: Barr's theory was that one or more bit flips resulted in the death of Task X, leaving the throttle valve stuck at near full throttle — and that none of the monitors or fail-safe devices detected the failure and shut down the throttle. Barr admitted they could not identify with certainty the specific ETCS software defects that resulted in UA, but told the jury that "to a reasonable degree of engineering certainty, it was more likely than not" that Task X death caused the accident.

- The plaintiff set out each element of a *prima facie* case in both negligence and strict liability by alleging: (1) Toyota owed the defendant a duty of care to make the ETCS reasonably safe; (2) Toyota breached that duty by negligently designing the ETCS with hardware and software defects; (3) that the negligently designed ETCS was the cause of the unintended acceleration of plaintiff's Camry; and (4) the plaintiff suffered injury and damage as a result of negligent design and defects in the Toyota ETCS.
- Plaintiff's expert witnesses were fully qualified and experienced in embedded software and software safety principles, and were appeared to be well coached by plaintiff's attorneys in product liability law. It appears that Toyota used an internal corporate expert who may not have had comparable skills. Barr's claim that the CPUs had no EDAC provisions was in conflict with NASA findings, but was apparently allowed to stand.
- Barr laid out a seemingly plausible set of events that could have resulted in the plaintiff's UA incident, albeit with a number of logical flaws and inconsistences [Ref. 11]. The causal link between the purported ETCS flaws and the Bookout incident were entirely speculative and theoretical, and could not be reproduced as described.
- The phrase "to a reasonable degree of scientific [or engineering] certainty" has no meaning in the scientific community, but is the wording recommended to expert witnesses in litigation. It has been criti-

cized as being confusing to a jury, and has been used to support both the preponderance of evidence and the reasonable certainty burden of proof [Ref. 14].

- In the Bookout case, the jury not only awarded compensatory damages, but had already decided to award punitive damages. This implies the jury found that Toyota was not only negligent in the ETCS design, but demonstrated a reckless disregard for the plaintiff under the higher burden of "clear and convincing evidence." Nothing in the record supports this conclusion, but Toyota settled the case before the amount of punitive damages was determined by the jury. This type of juror bias is well documented, and is one reason defendant companies will settle in the vast majority of cases [Ref. 15].

Safety Engineering Analysis of the Toyota ETCS

This section examines the safety engineering process undertaken by Toyota for the subject ETCS, and whether it conformed to known best practices. Also addressed is whether Toyota might have rebutted the purported design and process defects alleged by the plaintiff.

Determine Safety Goals, Requirements and Standards

Safety regulations for the Toyota ETCS included the Federal Motor Vehicle Safety Standard (FMVSS) for accelerator control systems [Ref 16]. The Main CPU used a RTOS based on the OSEK standard for distributed control units in vehicles, which is supported by Automotive Open System Architecture (AUTOSAR). Toyota claimed compliance with internal software coding standards that included roughly half of the voluntary MISRA C rules.

Critique: AUTOSAR provides architectural concepts that can enhance safety, such as memory partitioning, time determinism, program flow monitoring and communication stack-related features, etc. AUTOSAR also lists requirements for safety concerning data consistency, hardware memory protection features, data corruption detection and protection. Documented ETCS compliance with MISRA C guidelines and the safety provisions of the AUTOSAR standard would likely have rebutted most of the plaintiff's theory of bit flips and memory corruption. Subsequent to the Bookout litigation, a NHTSA study advised that the safety goals and requirements of MIL-STD-882E and ISO 26262 are applicable to the automotive industry [Ref. 17].

Hazard Identification and Analysis

It is assumed that Toyota identified UA as a high-severity hazard and considered the mitigations against this as "safety critical." It is unknown what hazard analysis techniques were performed by Toyota, but the material suggests at least FMEA was performed.

Critique: As part of its defense, Toyota contracted an independent consultant, Exponent Failure Analysis Associates, to investigate whether Toyota vehicles equipped with ETCS technology could accelerate without intentional command [Ref. 18]. Both Exponent and NASA employed standard safety analyses and tests to evaluate the safety of the ETCS, including FMEA, Fault Tree Analysis, software analysis and electromagnetic interference (EMI) testing. Neither Exponent nor NASA found single-point failures that could result in UA. NASA identified a scenario where two low-probability failures in the Pedal Assembly could mimic a valid accelerator pedal signal. In any event, NASA concluded that the ETCS was independent of the braking system. NHTSA demonstrated that a vehicle at high throttle could be stopped even under conditions of a depleted vacuum assist brake system.

Since both NASA and Exponent found the ETCS design to be reasonably safe and not a credible source of UAs, it is difficult to say that Toyota's original assessments were deficient, other than the allegation that the FMEA missed single points of failure. The record does not discuss hazard analyses performed on the ETCS during its design. Nor does it discuss whether the Camry brakes could stop a vehicle during a full power UA event, which seemed to be a point of disagreement at trial. A formal Safety Assessment Report (prior to release of the ETCS) documenting the safety mitigations for UA would have provided evidence of a good faith attempt to eliminate a *reasonably foreseeable risk of harm*.

Risk Assessment

It is assumed that Toyota conducted some form of risk assessment on the ETCS that included the well-known hazard of UA.

Critique: A risk assessment of UA during the design phase should have been documented by Toyota. The NASA Report recommended that controls for managing safety-critical functions, as currently applied to the railroad, aerospace, military and medical sectors, should be considered. Subsequent to the Bookout litigation, a NHTSA study advised that risk assessment methodologies of MIL-STD-882E and DO-178C are applicable to the automotive industry [Ref. 17].

Risk Reduction and Verification

Toyota designed many risk reduction measures into the ETCS, including a Main CPU with industry standard RTOS and memory EDAC provisions, as well as an independent Sub/Monitor CPU performing safety tasks of watchdog timer, monitoring critical parameters, diag-

> "Many companies are wary of the legal discovery process and have a document retention policy that purges files as quickly as possible. This may be short sighted. A manufacturer should have a robust product safety program that justifies safety design decisions, and a closed loop hazard tracking system that tracks hazards from requirements down to test and verification. It should be written with an eye toward future defect claims and discovery, documented in a Safety Assessment Report signed by management, and placed under configuration control."

nostics and transitioning the system to a fail-safe state in the event of a critical failure ("limp home" mode, engine power limiter and, finally, fuel cutoff). From calendar year 2005 to 2010, Toyota reported approximately 11 million hours in module-level software testing, and 35 million miles of system-level testing.

Critique: Exponent concluded that multiple "layers of protection" against resistive faults, component failures, software code and run time errors, bit flip/memory errors, latch-up and EMI mitigate the risk of UA. NASA reported that the ETCS "provides fail-safe modes to limit engine speed and engine power to a safe state to manage the risk" of UA. Taking into account the numerous safety design measures detailed in the Exponent and NASA reports, it is difficult to say that Toyota's ETCS safety architecture was deficient. However, lessons learned from the Bookout litigation and any other field and accident reports should be studied for consideration of upgrades or redesigns of the ETCS. The fact remains that Barr identified a number of deficiencies in the ETCS safety architecture. While the probability of the UA event occurring as alleged would seem to be extremely low, the ETCS should have been better protected against software errors and memory corruption.

Residual Risk Acceptance

It is unclear as to whether Toyota had an internal risk acceptance authority to determine whether the mishap risks had been reduced to an acceptable level — and either accept the residual mishap risk or require further risk reduction.

Critique: Assuming there was no internal risk acceptance authority, Toyota should have had a Product Safety Review Board (or equivalent) that included the design team, as well as safety engineers and legal representatives to provide risk acceptance of the ETCS design.

Suggestions to Mitigate Liability of Vehicle Manufacturers

The Bookout litigation provides many lessons for automotive manufacturers regarding the current acceptable level of safety of embedded software and complex control systems. Toyota seemingly fulfilled its duty to exercise a standard of care that complied with applicable regulations and produced an ETCS that was reasonably safe for its intended purpose. While the plaintiff's experts found theoretical flaws in the hardware and software design, they could not establish a direct causal link from the alleged defects to the UA event. They could, however, identify existing reasonable design alternatives that would likely have prevented that specific UA event. Unfortunately, this meets the burden of proof in civil litigation required for a plaintiff to claim that a theoretical defect caused a specific accident. The following suggestions are proffered to auto manufacturers for improvement in the areas of safety process and liability risk reduction in the current legal and regulatory environment.

Improve the Product/System Safety Process: Nearly any product safety program can be found wanting, particularly when a mishap involving that product has actually occurred (axiomatically making the risk "reasonably foreseeable"). The Toyota ETCS safety program is no exception, and a number of safety design features could have been improved, particularly in the areas of compliance with standards and software safety. Safety goals and safety requirements should have been specified. Full compliance with standards should have been demonstrated. The internal safety assessments should have been documented and released, along with evidence of risk mitigation verification, testing and closed loop hazard tracking. An independent safety assessment should have been standard policy. Finally, a corporate risk assessment and acceptance process for the risk of UA would have demonstrated a higher standard of care if conducted *prior* to the mishap.

Document the Safety Process: Many companies are wary of the legal discovery process and have a document retention policy that purges files as quickly as possible. This may be short sighted. A manufacturer should have a robust product safety program that justifies safety design decisions, and a closed loop hazard tracking system that tracks hazards from requirements down to test and verification. It should be written with an eye toward future defect claims and discovery, documented in a Safety Assessment Report signed by management, and placed under configuration control. This report should be completed prior to product release. If prepared correctly, the report can serve to rebut allegations of a defective or negligent safety process, as well as justify design decisions and counter theoretical design defects that may have little causal connection with an accident or mishap.

Voluntary Adoption of ISO 26262: The NASA Report recommended that controls for managing safety-critical functions be solicited from other industry domains. NHTSA identified several safety standards applicable to automotive electronic controls systems, including MIL-STD-882E, DO-178C and, most importantly, ISO 26262 [Ref. 19]. This functional safety standard is a requirement for E.U. automotive suppliers and is a voluntary standard available to U.S. manufacturers. It provides guidelines to accomplish a functional safety evaluation and to determine automotive safety integrity levels (ASIL). An ASIL designates a function's/item's essential safety requisites for attaining an *acceptable residual risk*. Hazard analysis and risk assessment is an important constituent of ISO 26262, with the objective to classify hazards of an item or function and develop safety goals (in terms of ASILs) to prevent or mitigate unacceptable risks and hazards. It is recommended that all automakers follow a functional safety process such as ISO 26262 and document compliance to the extent possible.

Mitigate Single Event Effects: Single Event Effects (SEE) on electronics resulting from ionizing radiation have long been a concern for high-altitude aircraft, but are increasingly being seen as a potential source of soft errors (e.g., bit flips) in high-density semiconductors used in ground systems, including motor vehicles. In fact, bit flips were cited as the most likely source of non-volatile memory corruption in the UA case study, and were considered in the NASA report. SEE error rates are difficult to estimate and may leave no trace of occurrence, making them an attractive mishap root cause that can neither be proved nor disproved. As such, it is recommended that vehicle manufacturers consider SEE analysis and mitigation on critical electronic components and assemblies. ISO 26262 provides some guidance on detecting and mitigating bit flips in embedded systems. Current best practice SEE mitigation techniques are available from the Federal Aviation Administration (FAA) [Ref. 20].

Adopt the ALARP Principle: In most of the E.U., it is a legal requirement that safety-related system residual risks be reduced to "As Low as Reasonably Practicable" (ALARP). "Reasonably practicable" is generally understood to undertake cost-benefit analysis, where the risk can be said to be reduced to a level that is ALARP if the cost of further reduction is "grossly disproportionate" to the reduction in risk gained. One legal interpretation of ALARP is that "good practice" in using safety codes and standards is automatically reasonably practicable. However, adequate good practice may not exist for the use of new technologies or complex systems. In this situation, the manufacturer needs to make a qualitative or quantitative judgement to support the ALARP argument [Ref. 21]. It is recommended to combine the functional safety approach of ISO 26262 to specify safety goals and ASILs with ALARP principles to determine whether an acceptable level of safety and a tolerable risk level have been achieved.

Promote the Involvement of NHTSA in Safety Certification: In the U.S., the FAA is authorized by law to set minimum standards for the design, materials, quality of work and performance of aircraft and their engines. The FAA regulations are comparable with the performance-oriented FMVSS promulgated by NHTSA, leaving the details of the design and development process to the manufacturer. Aircraft manufacturers must apply to the FAA for approval and certification to develop and build a new aircraft type. Each manufacturer applicant must present a certification plan that sets out the safety assurance processes it will use through development and production stages, including hardware and software quality and safety analyses and testing. The FAA must review and approve these plans, and test results before it grants certification for the aircraft or engine to be placed in service. Manufacturers are expected to implement safety assurance measures commensurate with the design assurance levels (DALs) for each safety-critical system, preferably by industry safety standards such as RTCA DO-178C.

Conclusion

In 2012, a committee sponsored by the Department of Transportation examined whether NHTSA should engage in comprehensive regulatory oversight of auto manufacturers, as occurs in the aviation industry. The committee concluded that NHTSA was not prepared to take on such a role [Ref. 22]. This government certification role should be revisited in light of the safety and liability challenges being faced by auto manufacturers — particularly in developing autonomous vehicles. The goal would be to provide a level of protection to domestic automakers such as that available in the E.U. [Ref. 23]. This approach of ISO standards, ALARP risk reduction and risk acceptance via government certification provides a potential path to improving safety and minimizing li-

ability for U.S. vehicle suppliers by preemption of federal law [Ref. 24][3].

About the Author

Martin S. Chizek, P.E., C.S.P. is a product safety officer at Lockheed Martin in Orlando, Florida. He holds a Bachelor of Science in Engineering from Wichita State University, a Master of Science in Systems Engineering and Management from University of Southern California, and a Juris Doctor from University of Memphis.

[3] "We hold that Plaintiffs' failure to warn claim is preempted by the Federal Aviation Act, which together with the regulations promulgated by the Federal Aviation Administration, exclusively governs the entire field of aviation safety."

References

1. "Military Standard 882E," U.S. Department of Defense, May 11, 2012, http://www.system-safety.org/Documents/MIL-STD-882E.pdf.
2. "DO-178C, Software Considerations in Airborne Systems and Equipment Certification," Radio Technical Commission for Aeronautics, Inc. (RTCA) and the European Organization for Civil Aviation Equipment (EUROCAE), 2011.
3. Robinson v. Brandtjen & Kluge, Inc., 500 F.3d 691 (8th Cir. 2007).
4. Restatement (Second) of Torts § 402A, "Special Liability of a Seller of Product for Physical Harm to User or Consumer," American Law Institute, 1965.
5. "Restatement (Third) of Torts, Product Liability," American Law Institute, 1997.
6. Knitz v. Minster Machine Company, 69 Ohio St. 2d 460, 432 N.E.2d 814 (1982).
7. Swanson, Derek H. and Dr. Lin Wei. "U.S. Automotive Products Liability Law, A Corporate Approach to Preventive Management, Risk Reduction, and Case Coordination for Chinese Automakers," McGuireWoods, LLC., 2011.
8. NASA Engineering and Safety Center (NESC). "Assessment TI-10-00618, NASA Engineering and Safety Center Technical Assessment Report, NHTSA Toyota Unintended Acceleration Investigation," January 2011.
9. U.S. Department of Transportation. "NHTSA Technical Assessment of Toyota Electronic Throttle Control System," February 2011.
10. Case No. CJ-2008-7969, Jean Bookout, et al. V. Toyota Motor Corporation, District Court of Oklahoma County, Oklahoma, October 2013, http://www.safetyresearch.net/Library/BarrSlides_FINAL_SCRUBBED.pdf.
11. David M. Cummings. "Embedded Software Under the Courtroom Microscope — A Case Study of the Toyota Unintended Acceleration Trial", *IEEE Technology and Society Magazine*, December 2016.
12. Chris Isadore. "Toyota Settles Acceleration Case After $3M Jury Verdict", *CNN Money*, October 25, 2013.
13. "Development Guidelines for Vehicle Based Software," The Motor Industry Software Reliability Association (MISRA), November 1994.
14. National Commission on Forensic Science. "Testimony Using the Term "Reasonable Scientific Certainty," National Institute of Science and Technology (NIST), 2014.
15. Viscusi, W. Kip. "Does Product Liability Make Us Safer?" *Regulation*, Vol. 35, No. 1; Vanderbilt Public Law Research Paper No. 12-20, Spring 2012.
16. 49 C.F.R. Part 571, Federal Motor Vehicle Safety Standards.
17. "Assessment of Safety Standards for Automotive Electronic Control Systems," Report No. DOT HS 812 285, National Highway Traffic Safety Administration, June 2016.
18. "Analysis of Toyota ETCS-i System Hardware and Software," Exponent Failure Analysis Associates, September 2012.
19. "ISO Standard 26262: Road Vehicles — Functional Safety," International Organization for Standardization (ISO), November 2011.
20. "DOT/FAA/TC-15/62, Single Event Effects Mitigation Techniques Report," Federal Aviation Administration, February 2016.
21. "Reducing Risks, Protecting People," Health and Safety Executive, 2001.
22. "Transportation Research Board Special Report 308, The Safety Challenge and Promise of Automotive Electronics: Insights from Unintended Acceleration," National Academies Press, 2012.
23. "Directive 2007/46/EC, Framework for the Approval of Motor Vehicles," European Commission, September 2007.
24. Montalvo v. Spirit Airlines, 508 F.3d 464 (9th Cir. 2007).