

INVITED: Extensibility in Automotive Security: Current Practice and Challenges

Sandip Ray¹, Wen Chen¹, Jayanta Bhadra¹, Mohammad Abdullah Al Faruque²

¹NXP Semiconductors, Austin, TX 78735. USA. firstname.lastname@nxp.com

² University of California at Irvine, Irvine, CA 92697. USA. alfaruqu@uci.edu

ABSTRACT

A modern automotive design contains over a hundred microprocessors, several cyber-physical modules, connectivity to a variety of networks, and several hundred megabytes of software. The future is anticipated to see an even sharper rise in complexity of this electronics, with the imminence of driverless vehicles, the potential of connected automobiles within a few years, and work towards seamless integration of automobiles with smart cities and infrastructure systems. Security is a fundamental challenge in the design of automotive systems. Unfortunately, security considerations in automotive systems are complicated by two factors: (1) need for real-time mitigation against in-field threats; and (2) in-field configurability and extensibility of security features. This paper examines the trade-offs between security countermeasures, real-time requirements, and in-field configurability needs for modern automotive systems. We discuss the current state of the practice in automotive security architecture, as well as gaps and challenges that need to be addressed for a viable security solution in future.

Keywords

functional safety, side channel, V2X, vehicle security

1. INTRODUCTION

Recent years have seen a rapid increase in complexity of electronics and software components in our automobiles. A modern car can contain over 200 electronic control units (ECU), several in-vehicle communication networks, and several hundred megabytes of software. Indeed, electronics and software are now viewed as the key market differentiators for automotive systems and also account for more than 50% of the design overhead [22]. As we move in the era of increasing automation in vehicles, this complexity is only slated to grow more sharply. An obvious consequence of this complexity is the rise in defects due to electronic and software errors, resulting in recalls as well as demonstrations showing

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DAC '17 June 18–22, 2017, Austin, TX, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4927-7/17/06...\$15.00

DOI: <http://dx.doi.org/10.1145/3061639.3072952>

how a vehicle can be hacked or controlled from outside by a malicious agents [16, 15].

All this is new to the culture and tradition of the automotive industry, that has historically focused on a much slower trajectory of complexity growth than electronic systems together with more careful attention to quality, safety, and robustness. The traditional distinction between automotive and electronic systems was put very succinctly by the following “tongue-in-cheek” quotation from Paar [22]:

If vehicles were developed in the same manner as telecommunications, then an average car would reach top speeds of 10⁹km/h at 400M HP and the car would be hacked four times a year.

Now that vehicles *are* being developed as electronic systems, one must account for the corresponding different trajectory of automation in a modern automobile, and additionally reconcile this growth (and its accompanying system vulnerabilities) with the expectations of robustness, safety, and security that the customers have traditionally expected out of a car.

A key requirement for robust architecture is *extensibility*, *i.e.*, ensuring that it can be easily extended to address future requirements. A key challenge to extensibility in automotive system architectures stems from the long in-field life of the system. In contrast to other common consumer electronic systems (*e.g.*, a mobile phone which has a typical field life of about a year), a car has a very long field life, lasting more than a decade. Within the course of this time, security requirements, solutions, and even attacks may change significantly. Consequently, one must design an automotive system with significant in-field configurability. Furthermore, the architecture must build in a flow for in-field updates which itself must be upgradable as technology advances.

In this paper, we discuss some challenges in architecting extensible security architecture for automotive systems. Some of the challenges are unique to automotive systems; some stem from the fact that an automotive system, particularly in the impending future world featuring self-driving, autonomous cars, is anticipated to be a ubiquitous part of a seamlessly connected Internet of Things infrastructure. We discuss the current state of the practice in automotive architecture design, approaches taken today to ensure extensibility, their limitations for future requirements, and the trade-offs that must be accounted for in addressing them.

The remainder of the paper is organized as follows. Section 2 discusses current trends in automotive system design, and underlines the critical role of electronic and software components. Section 3 introduces the robustness challenges

in automotives, involving interplay between safety and security components. We go a bit deeper into security in Section 4, discussing in some detail security attack models and modes of attack. Sections 5 and 6 discuss extensibility issues from the perspective of security, including the requirements that drive extensible architectures and the trade-offs and challenges one faces to realize one. Section 7 discusses the current state of practice in extensible architectures and outlines some recent research directions. We conclude in Section 8.

2. TRENDS IN AUTOMOTIVE SYSTEMS

Starting from the 1990s, the design complexity of automotive systems has been dominated by electronic parts, with more focus on software components in the last decade. The electronic and software components in an automotive design (which we will loosely refer to as “electronics”) can be classified into three major components:

Infotainment Components: The goal of infotainment components has been traditionally to provide real-time road and traffic information as well as entertainment in the vehicle. Infotainment has also historically been the motivator that enabled seamless connectivity to wireless networks. Most modern automotive systems include several microcontroller units devoted to infotainment, in addition to broadcast radio, connectivity to the user’s mobile devices via bluetooth, USB, and other protocols, connectivity to Internet to enable communication of roadside emergency, etc. With the emergence of more and more autonomy in cars, infotainment in future vehicles may include connection to several networks, *e.g.*, roadside cafeterias, targeted applications provided by the automotive companies, repair shops, and diverse third-party applications.

Driver Assistance: The automated driver assistance systems (ADAS) include electronic and software components targeted towards assisting in, enhancing, and eventually replacing the functions of the human driver. ADAS includes a variety of sensors, camera, radar, and LIDAR devices to identify the surroundings, determine wheel speed and angular momentum, etc. Typically, sensor data is accumulated into a Sensor Fusion module that performs analytics to identify various road conditions (*e.g.*, pot-holes, pedestrians, other vehicles, etc.). Inference from the sensor fusion may activate various (electro-mechanical) vehicle controls including automated braking, steering control, etc. Note that ADAS typically has connections with the infotainment components as well, *e.g.*, sensor fusion can make use of maps and terrain information provided by the GPS.

Environmental Safety: A key requirement for vehicles today, — particularly as legal requirements are being crafted for autonomous ones — is to operate in an environmentally responsible manner. For example, the European Union mandates reduction of emissions by 20% for vehicles on the road by 2020.

3. THE AUTOMOTIVE ROBUSTNESS PROBLEM

What do we mean by a robust and trustworthy automotive system? Roughly, *automotive robustness* refers to the following four requirements.

Functional Safety: This is the requirement that the automobile should not harm other agents (human, other vehicles, etc.) due to hazards caused by malfunctioning behavior of electrical/electronic (E/E) systems. Requirements from functional safety are embodied by the ISO 26262 standard, which is an adaptation of the functional safety standard IEC 61508 for automotive E/E systems. Functional safety requires that the system should be able to mitigate hazardous events that could occur during the execution of the automotive E/E systems. Five risk levels called Automotive Safety Integrity Levels (ASIL) are defined for the potential hazards, ranging from the highest level of hazard (ASIL D) to the non-hazardous level (QM). The ASIL of a hazardous event is defined in terms of the severity of the potential injury it can cause, and the likelihood of the hazard happening and the likelihood that the driver can act to prevent injury. Functional safety requires the system to be fault tolerant, meaning that the system can operate correctly under certain level of systematic faults and random hardware faults. Single Point of Failure (SPF), where a single fault stops the system from working, is unacceptable for automotive E/E systems.

Security: Security refers to the requirement that the electronic components of the car must be resilient against system hacks. An informal way to distinguish security from functional safety is to note that while functional safety requires that the car does not harm others, security requires that other agents should not be able to harm the car. We discuss security challenges in more detail in the subsequent sections.

Device Reliability: Device reliability, or component reliability, refers to robustness of electromechanical and mechanical components in the vehicle. Approaches to ensure device reliability include sensors and analytics software for providing early warning against component wear-outs, mechanisms to ensure slow and gradual degradation, etc.

Road Safety: This refers to the key goal of minimizing accidents and road hazards stemming from human errors. Note that human errors account for more than 90% of road accidents today. As discussed in Section 2, a primary objective of ADAS is to diminish incidents caused by human errors. Eventually, as we move towards full automation, the goal is to eliminate the road safety component from robustness criteria by eliminating the human driver from the loop.

Clearly, the four components above are inter-related. For example, an external hack can cause the system to fail in a way that harms other agents, reducing functional safety to a security issue. Correspondingly, a malfunction in ADAS (because of functional safety, security, or component failure) may result in a road safety incident, *e.g.*, by failing to warn a driver against an approaching pedestrian.

4. AUTOMOTIVE SECURITY BASICS

As discussed above, security is only one of several components in the automotive robustness problem. Security itself, however, is a complex and pervasive problem with a wide range of challenges. Understanding security challenges for most requires comprehension of two distinct but interrelated components, *e.g.*, attack models, and attack modes. Attack models refer to the objectives of the attacker. The objectives include subversion of confidentiality, integrity, or availability

requirements of the system. Below, we give a brief, high-level overview of the various attack models in automotive systems. Attack modes refer to the actual operations that an attacker can perform to subvert these objectives. For an automotive, attack modes typically involve ways to hijack the car’s electronic mechanisms to gain unauthorized access to core driving functions or compromise the privacy (*e.g.*, location, infotainment content, etc.) of the vehicle in an unauthorized manner. In addition, vehicle security includes a third component, *e.g.*, gaining unauthorized physical access to the car (*e.g.*, for theft). This third component includes some orthogonal challenges, and we will refer to this third component as *physical access security*.

4.1 Attack Models

Below, we present physical attack models on the automotive system and classify them according to the security requirements that they breach: **confidentiality**, **integrity**, and **availability**. Attackers may have different motivations (*e.g.*, financial or political gain, injury, etc.), budgets and resources, and levels of knowledge. Furthermore, there are different vulnerabilities that an attacker can utilize for their attack, including sensors, hardware, and software.

Confidentiality: In the confidentiality attack model, the attacker aims to derive some critical information (*e.g.*, location, habits, etc.) either about the user or about the system. Traditionally, confidentiality attackers have breached the software of the system to derive critical information. Although more difficult, an attacker may also invade the physical domain to attain this information because it is less detectable and cost-efficient (*e.g.*, requires only sensors and remote connection). Furthermore, with sufficient quantity and quality information about the physical components and channels, an attacker can also determine this critical information without needing to be physically near system. By having sensors placed near the components (*e.g.*, GPS, TPMS, battery) of the system, the attacker may extract critical information from GPS, battery usage, temperature, velocity, acceleration, tire pressure values [11], and more. With these values, an attacker can derive the physical location and surroundings of the system as well as the personal habits of the driver. Depending on the type of the vehicle, some values may have more weight than others, *e.g.*, localization features for autonomous, battery features for electrical vehicle (EV), etc.

Integrity: In the integrity attack model, an attacker intends to modify legitimate values to deceive either the user or the system. This includes approaches to corrupt sensitive data, or tamper with communications (either among components of a vehicle, *e.g.*, through in-vehicle networks, or messages between a vehicle and the external world). Furthermore, an attacker may perform straightforward physical damage or tampering with the physical domain components and/or channels.

Availability: In the availability attack model, an attacker desires to deny the user or system of a service. By modifying some components and/or channels in the physical domain, the attacker can create a denial of service (DoS) attack. Some examples for these types of attacks are as follows. An attacker can spoof the incoming GPS signal of the vehicle [9, 18], the LIDAR signals from autonomous vehicles [7], the battery values from the battery system in EVs

(via battery counterfeit/ swapping/modification [21, 23] or battery management system hacking [14]), the tire pressure sensor values (from the Tire Pressure Monitoring System sensor network [11]), and even potentially the MEM-based accelerometers around the vehicle by using sound [13].

4.2 Security Attack Modes

An autonomous car must communicate with a number of agents during to perform various ADAS and infotainment functions. This exposes the electronic components in the vehicle to malicious attacks. Recent ethical hack¹ demonstrations [15] make such hacks a serious and critical concerns as we move towards increasingly higher degree of autonomy. Note from above that this includes both security and privacy concerns. However, the two concerns are somewhat different. The following two hypothetical scenarios illustrate the concerns.

- **Security Scenario:** Suppose an autonomous car receives a message from another vehicle or the highway system (*e.g.*, an approaching vehicle warning). It must be able to trust the identity of the sender and the integrity of the message. Failure to do so can easily enable a malicious agent to cause significant traffic disruption or accident.
- **Privacy Scenario:** Suppose an autonomous car sends a message to other vehicles or roadside assistance, *e.g.* to provide a hazard warning, request for traffic information, etc. This exposes the car to potential privacy violations: malicious agents listening in to the communication can track its location, direction, and identity. To ensure privacy, it is critical that such communications are anonymized.

The two above scenarios demonstrate an interesting conundrum between security and privacy: trusting in-field communications requires the ability to verify authenticity of the sender which may be in conflict with the sender’s privacy concerns. Architectural solutions to address the issue must comprehend the subtle trade-offs between authentication and anonymization requirements in order to be viable.

To better understand architectural requirements for (functional) security assurance in automotive systems, we now consider the high-level sources of vulnerability. Roughly, there are two critical challenges.

Side-channel Leakage: Side-channel problems arise in automotive systems since the attacker can have physical access to the device. This permits the attacker to have access sensitive information through various side-channel sources. Of course, side channel leakage is nothing new: it has been studied in the context of hardware security for decades [12]. However, the large number of electronic components in a car, as well as the complexity of their coordination, give rise to complex side-channels that are difficult to anticipate in advance. For example, in addition to traditional side-channel leakage through voltage, current, and thermal profiles, sensitive cryptographic and DRM keys in an automotive system

¹An *ethical hack*, also referred to as a *penetration test* or *intrusion test*, is an attack performed by scientists and engineers without malicious intent to demonstrate a possible vulnerability of an electronic device to a real hack in-field by mimicking the capabilities of a malicious agent.

can be leaked through a study of emission profile, latency of reaction in a specific in-field situation (*e.g.*, pot-hole or pedestrian avoidance), etc. We discuss these issues in more detail below.

In-field Communication Requirements: Since automotive systems must operate in-field for a long time, they must include facilities for in-field, over-the-air (OTA) updates to software, firmware, or even hardware configurations. Furthermore, automated operation depends critically on the ability of the vehicle to communicate in real time with other vehicles as well as the roadway systems. Since billions of electronic devices are connected with these networks, one can assume that it includes several (hundreds of thousands, if not millions) of malicious or compromised systems. As explained in the security and privacy scenarios above, such communication comes with risk of exposure of sensitive, security-critical information in the car to unauthenticated or malicious agents. Again, note that the challenge is not unique to automotive systems: any connected computing device that needs to transmit sensitive or personalized information is vulnerable to this problem. However, the problem is acute for autonomous automotive systems in particular because its core functionality critically depends on real-time communication with other agents.

To illustrate how the side-channel and communication challenges can together enable powerful attacks, consider the following scenario. OTA updates from the manufacturer require access to secure cryptographic keys: if an attacker can have access to these keys then they can use it to install arbitrary — potentially malicious — software through the OTA updates. Now consider an adversary who has physical access to one of these vehicles. It is then possible to get access to the cryptographic keys *for that vehicle* via side-channel leakage. Subsequently, the adversary can use the key so extracted to perform malicious software updates to *other similar vehicles*. The key reason for this ability is that many electronic components are produced *en masse* with the same configuration of keys and other information. While this can be a simple oversight in production, in practice it is difficult to ensure (and manage) unique configuration of keys for each device. Consequently, one compromised ECU can lead potentially severe security compromise of a whole class.

We end this subsection with the discussion promised above on unique side-channel challenges in automotive systems. The challenges here arise from the fact that an automotive is a cyber-physical system (CPS), and a cyber-physical system is vulnerable to side channels based on both “physical” and “cyber” characteristics. For example, it has been recently shown that in a CPS (*e.g.*, in a 3D printer), an attacker can take advantage of emissions (*e.g.*, acoustic) from the physical domain in a way to extract valuable information meant to be secured [2]. In general, with a creative mind, an attacker can take advantage of emissions from the physical domain in a way to extract information meant to be secured [1, 2]. However, kinetic-cyber attacks are not limited there. In fact, there are even attacks that aim to modify the physical domain to manipulate the system’s functionality.

4.3 Access Security

Another component of an automotive security is *physical accessibility and authenticity*. The goal of this component is to ensure that only the rightful owner is able to get physical access to the vehicle at any time. Note that if an attacker can

get physical access to the vehicle, then in addition to stealing the vehicle they may also be able to tamper its mechanical, electro-mechanical, or electronic systems to further breach the confidentiality, integrity, and availability requirements.

Vehicle manufacturers have been using software-assisted authentication (physical key with RFID, keyless entry with RFID, and passive keyless entry and start) for physical access because physical keys are no longer useful [6]. However, researchers have discovered several vulnerabilities with these new authentication schemes. For example, in [8], researchers demonstrated that a keyless fob can be hacked by relaying the signal from a wireless transceiver (at either low, high, or ultra high frequency) via cable or over the air. In another example, an attacker could start a vehicle’s engine by exploiting the vulnerabilities of certain Digital Signature Transponders to reverse engineering, key cracking, and simulation attacks [5].

5. EXTENSIBILITY DRIVERS

The diversity of security challenges discussed clearly suggest the requirement of a disciplined, robust security architecture for the future. One crucial component of such an architecture, however, is that be *extensible*, *i.e.*, not only should it facilitate mitigation of today’s security threats, but it must also enable easy in-field configurability to facilitate adherence to future requirements. There are several factors that drive the need for extensibility. In this section, we enumerate some of these factors.

Long In-field Lifetime: We briefly touched upon this point in the introduction. Compared to most electronic systems, a car has a much longer in-field life-span, going more than a decade. Looking back at the change in security requirements and implications over the past 15 years, we can imagine that the requirements of future within this life-time may be radically different from today. Second, even if the usage model does not change, such a long life-time is well beyond the horizon of trustworthiness for current trust protection mechanisms, *e.g.*, hardware implementation for a modern cryptographic encryption algorithm has an anticipated assurance time-frame of about 5 to 7 years. Third, protection requirements themselves may change during the lifetime of the system, possibly in response to new attacks discovered when the device is on field.

Dynamic Trade-offs between Security, Smartness, Communication: An autonomous car in operation must make real-time decision on trade-offs between security, energy, and smartness. For example, a car driving on a desolate, straight highway requires less data analytics for pot-hole or pedestrian detection than when driving in a busy city; this enables the car to adjust its communication bandwidth to the cloud in real time. Addressing such dynamic, real-time adjustments requires that the underlying architecture provide generic interfaces for communication (as well as clear definition of various communication, smartness, and security modes). Furthermore, with increasing advancement of detection and analytics software, it will be possible to support increasingly smarter trade-offs in future. Consequently, it is critical for the underlying architecture to be extensible to accommodate the possibility.

Communication Standardization Needs: An automotive today must continuously communicate with other vehicles, with the highway infrastructure, and with billions of

other electronic components. Indeed, it is this communication that characterizes the Internet-of-Things regime [19]. For automotive systems, this communication — referred to as “V2X” — is a cornerstone of autonomous functionality. On the other hand, the infrastructure architecture, communication protocols, and communication functionality, are constantly evolving. Even in cases where the protocols themselves are standardized, communication patterns can govern various trade-offs between security, performance, and network bandwidth utilization. To address all of this, it is critical that it be possible for the V2X infrastructure to be extended in a disciplined manner to maintain coherence with evolving interfaces and communication patterns with the rest of the ecosystem.

Bulk Production Needs: Electronic components for automotive systems are created in “bulk”, and typically reconfigured and tuned for various in-field needs. Enabling this approach requires that these components be developed as generic, reconfigurable (and hence extensible) hardware modules rather than as custom hardware.

Verification Needs: The diversity of security challenges in automotive suggests a diverse range of security mechanisms. A resultant requirement is to verify that the automotive system satisfies the security requirements under a plethora of use case scenarios. For example, it is necessary to verify that the V2X communication remains secure regardless of how many vehicles and RSUs are in proximity. Due to the resource constraints in verification and time-to-market requirements, it is not feasible to verify all possible configurations. Therefore, it requires the verification results under certain configurations can be extended to other configurations. Ultimately, this translates to the requirement in architectural extensibility.

6. CHALLENGES TO EXTENSIBILITY

While extensive secure mechanisms have been deployed to protect current automotive systems from existing attacks, providing true in-field extensibility for future generations of automotive systems is a challenging problem. It is crucial to understand the trade-offs and complexities induced by extensibility needs. In this section we enumerate some of these challenges. An architect developing an automotive security architecture must balance these constraints against the benefits of extensibility.

Optimization Needs: A crucial trade-off blocking extensibility is the need for custom optimization for specific performance or security needs. Conflicts between extensibility and optimization are of course, well-known and not germane to automobiles. However, optimization needs, particularly in real-time requirements, are crucial for automotive systems making the trade-off more acute.

Verification Needs: As mentioned above, verification is a crucial driving factors for extensibility. Unfortunately, there is also trade-off between verification and extensibility. In particular, extensibility typically involves developing an architecture with more behaviors and configurations than necessary for current use cases. This puts the burden on verification to ensure correctness of these additional configuration, — a task made more challenging because they are essentially “reserved for future use” with no unambiguous functionality requirement for the present. Even so, verification needs to account for them particularly in the context of

security, since such unused configurations and behaviors are typical targets of security vulnerabilities.

Time-to-Market: A final, critical challenge to extensibility is time-to-market constraint. Extensible architectures typically *reduce* time-to-market in future products. However, they have longer latency of development at first deployment. Unfortunately, in modern, highly competitive business environment, it is often difficult to tolerate the short-term latency and enable extensibility for the long term.

7. STATE OF THE PRACTICE: A SECURITY ASSURANCE ARCHITECTURE

There has been significant work in industrial practice on developing disciplined architectures for security requirements for automotive systems. Here we discuss one approach, referred to as the 4+1-layer security assurance architecture [17]. In this architecture, security concerns are defined hierarchically in four layers, as defined below. Note that more generic in-field extensibility may be needed to make sure that such architectures can be readily extended to meet the security requirements of future automotive systems.

Secure Interfaces: This layer includes securing communication of the car with the external world. Components of Layer 1 security architecture includes V2X, telematics, etc. In the setting of V2X, vehicles and road-side units (RSU) can broadcast messages to any vehicles equipped with V2X receivers within a range. IEEE 1609.2 has been proposed to ensure privacy, authenticity and non-repudiation of the V2X message communication using a set of cryptographic techniques such as encryption, digital signature and certificates. For communication between vehicles and the cloud, existing Internet security technologies such as HTTPS and TLS can be leveraged for protection.

Secure Gateway: This layer acts as a firewall between the external interfaces and the safety-critical in-vehicular networks (IVN). It monitors and controls the traffic coming into the trusted IVNs from the outside world and filters out potentially hazardous transactions. The secure gateway also plays a role in routing traffic from one IVN to another. In case one IVN is compromised, the gateway can isolate the compromised components and prevent the attack from propagating to other IVNs.

Secure Networks: This layer includes securing the IVNs and ensuring that attackers cannot take advantage of the compromise of one IVN to invade other IVNs. Unfortunately, most commonly used IVN protocols such as LIN, CAN and FlexRay lack security mechanisms. Currently, the central gateway is heavily relied on for providing logical and physical isolation between IVNs. Automotive Ethernet, the next-generation IVN protocol, is supposed to provide more intrusion detection capabilities and stricter separation.

Secure Processing: This layer involves securing the MCU and MPU units. As many complex functionalities are implemented as software/firmware to maximize design productivity and in-field extensibility, it is critical to ensure authenticity and integrity of the firmware running on the MCU/MPU units. These units are equipped with hardware implementation of the Secure Hardware Extension (SHE) specification to accelerate the cryptographic operations for enabling such authentication. Virtualization is employed to realize process isolation to prevent one compromised software stack from

being exploited to attack other software stacks. Tamper detection and resistance mechanisms are often implemented to protect MCU/MPUs from voltage/clock manipulation.

In addition to the above four layers, the “+1” layer provides physical vehicle protection through anti-theft immobilizer and smart car access functions. Innovations in this area include new features like remote lock and unlock, passive start, remote vehicle monitoring, and car access using a smart phone or smart key device.

Research Directions: There has also been other recent research with the goal to develop disciplined automotive security architectures. For example, a flexible security architecture has been proposed in recent work [20, 3, 4] that enables centralized specification of security requirements for MCUs and MPUs; this work specifically targets flexibility to enable in-field upgrades. There have also been efforts at developing disciplined methodology towards defining functional safety and security goals in an extensible manner [10]. In spite of these advances, significant work remains to address the problem of extensibility for future systems in a way that accounts for the diversity of constraints discussed above.

8. CONCLUSION

We have discussed security and trustworthiness challenges in automotive systems, the need for extensibility, and the constraints and considerations involved in achieving it. We also provided a flavor of the state of the practice in automotive security architecture today, and how they are limited in the degree of genericity and extensibility for the expected requirements of tomorrow’s automotive systems. While we have discussed some directions of research, we are currently only scratching the surface of this challenging problem. A comprehensive solution will require a re-thinking of the architecture, and comprehending and accounting for the trade-offs necessary among various stakeholders’ interests.

9. REFERENCES

- [1] D. Agrawal, B. Archambeault, J. R. Rao, and P. Rohatgi. The em side—channel(s). In B. S. Kaliski, ç. K. Koç, and C. Paar, editors, *CHES 2002*, 2002.
- [2] M. A. Al Faruque, S. R. Chhetri, A. Canedo, and J. Wan. Acoustic side-channel attacks on additive manufacturing systems. In *7th International Conference on Cyber-Physical Systems*, pages 19:1–19:10. IEEE Press, 2016.
- [3] A. Basak, S. Bhunia, and S. Ray. A Flexible Architecture for Systematic Implementation of SoC Security Policies. In *ICCAD*, 2015.
- [4] A. Basak, S. Bhunia, and S. Ray. Exploiting design-for-debug for flexible SoC security architecture. In *DAC*, 2016.
- [5] S. Bono, M. Green, A. Stubblefield, A. Juels, A. D. Rubin, and M. Szydlo. Security analysis of a cryptographically-enabled rfid device. In *USENIX Security*, volume 5, pages 1–16, 2005.
- [6] D. Clare, S. Fry, H. Handschuh, H. Patil, C. Poulin, A. Wasicek, R. Wood, D. Brown, G. Cooper, I. Gilvary, D. Grawrock, A. Rajan, A. Tatourian, R. Venugopalan, C. Vishik, D. Wheelere, and M. Zhao. Automotive Security Best Practices: Recommendations for Security and Privacy in the Era of the Next-generation Car. Technical report, McAfee and Intel.
- [7] J. Condliffe. A \$60 hack can fool the lidar sensors used on most self-driving cars, 2015.
- [8] A. Francillon, B. Danev, and S. Capkun. Relay attacks on passive keyless entry and start systems in modern cars. In *18TH Annual Network and Distributed System Security Symposium*, 2011.
- [9] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, B. W. Oâ€™Hanlon, and P. M. Kintner Jr. Assessing the spoofing threat: Development of a portable gps civilian spoofer. In *Proceedings of the ION GNSS international technical meeting of the satellite division*, volume 55, page 56, 2008.
- [10] Intel. Car of the Future – Automotive Safety and Security Trends. <http://www.intel.com/content/www/us/en/automotive/automotive-security-best-practices-white-paper.html>.
- [11] R. M. Ishtiaq Roufa, H. Mustafaa, S. O. Travis Taylora, W. Xua, M. Gruteserb, W. Trappeb, and I. Seskarb. Security and privacy vulnerabilities of in-car wireless networks: A tire pressure monitoring system case study. *19th USENIX Security Symposium, Washington DC*, pages 11–13, 2010.
- [12] P. C. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In *Proceedings of the 19th Annual International Cryptology Conference on Advances in Cryptology*, pages 388–397, 1999.
- [13] M. K. McGee. Study: Some mobile devices can be hacked using sound waves, 2017.
- [14] C. Miller. Battery firmware hacking: Inside the innards of a smart battery. *Black Hat USA*, 2011.
- [15] C. Miller and C. Valasek. Remote Exploitation of an Unaltered Passenger Vehicle. In *BlackHat USA*, 2015.
- [16] National Highway Traffic Safety Association. Motor Vehicles Recall. See URL: <https://www.recalls.gov/nhtsa.html>.
- [17] NXP Semiconductors. <http://www.nxp.com/assets/documents/data/en/white-papers/MULTI-LAYER-VEHICLE-SECURITY-WP.pdf>.
- [18] M. Psiaki and T. Humphreys. Protecting gps from spoofers is critical to the future of navigation, 2016.
- [19] S. Ray, Y. Jin, and A. Raychowdhury. The Changing Computing Paradigm with Internet of Things: A Tutorial Introduction. *IEEE Design & Test of Computers*, 33(2):76–96, 2016.
- [20] S. Ray, J. Yang, A. Basak, and S. Bhunia. Correctness and Security at Odds: Post-silicon Validation of Modern SoC Designs. In *DAC*, 2015.
- [21] F. Sagstetter, M. Lukasiewicz, S. Steinhorst, M. Wolf, A. Bouard, W. R. Harris, S. Jha, T. Peyrin, A. Poschmann, and S. Chakraborty. Security challenges in automotive hardware/software architecture design. In *DATE*, pages 458–463, 2013.
- [22] A. Weimerkirsch. Automotive and Industrial Data Security. In *Cybersecurity and Cyber-physical Systems Workshop*, 2012.
- [23] S. Weintraub. Gogoro scooter and battery swap distribution model get smarter and spread to new cities, 2016.