On the Feasibility of Carrying Out Live Real-Time Forensics for Modern Intelligent Vehicles

Saif Al-Kuwari[1,2] and Stephen D. Wolthusen[1,3]

[1] Information Security Group, Department of Mathematics, Royal Holloway,
University of London, Egham Hill, Egham TW20 0EX, United Kingdom
[2] Information Technology Center, Department of Information and Research,
Ministry of Foreign Affairs, P.O. Box 22711, Doha, Qatar
[3] Norwegian Information Security Laboratory, Gjøvik University College,
P.O. Box 191, N-2802 Gjøvik, Norway

Summary. Modern vehicular systems exhibit a number of networked electronic components ranging from sensors and actuators to dedicated vehicular subsystems. These components/systems, and the fact that they are interconnected, raise questions as to whether they are suitable for digital forensic investigations. We found that this is indeed the case especially when the data produced by such components are properly obtained and fused (such as fusing location with audio/video data). In this paper we therefore investigate the relevant advanced automotive electronic components and their respective network configurations and functions with particular emphasis on the suitability for *live* (real time) forensic investigations and surveillance based on augmented software and/or hardware configurations related to passenger behaviour analysis. To this end, we describe subsystems from which sensor data can be obtained directly or with suitable modifications; we also discuss different automotive network and bus structures, and then proceed by describing several scenarios for the application of such behavioural analysis.

Keywords: Live Vehicular Forensics, Surveillance, Crime Investigation.

1 Introduction

Although high-speed local area networks connecting the various vehicular subsystems have been used, e.g. in the U.S. M1A2 main battle tank[1], complex wiring harnesses is increasingly being replaced by bus systems in smaller vehicles. This means that functions that had previously been controlled by mechanical/hydraulic components are now electronic-based, giving raise to the X-by-Wire technology [1], potentially turning the vehicle into a collection of embedded interconnected Electronic Control Unites (ECU). However, much of the recent increase in complexity has arisen from comfort, driving aid, communication, and

[1] Personal communication, Col. J. James (USA, retd.).

X. Lai et al. (Eds.): E-Forensics 2010, LNICST 56, pp. 207–223, 2011.

entertainment systems. We argue that these systems provide a powerful but as-yet under-utilised resource for criminal and intelligence investigations. Although dedicated surveillance devices can be installed at the in-vehicle system, these are neither convenient nor economical. On the other hand, the mechanisms proposed here can be implemented purely in *software* and suitably obfuscated. Moreover, some advanced automotive sensors may also provide redundant measurements that are not being fully used by the corresponding function, such as vision-based sensors used for object detection where images/video from the sensor's measurements are inspected to detect the presence of objects or obstacles. With appropriate modifications to the vehicular electronic systems, this (redundant) sensor information can then be used in forensics investigation. However, the fact that components are interconnected by bus systems implies that only central nodes, such as navigation and entertainment systems, will need to be modified and can themselves collect sensor data either passively or acquire data as needed. We also note the need for awareness of such manipulations in counter-forensic activity, particularly as external vehicular network connectivity is becoming more prevalent, increasing the risk, e.g., of industrial espionage.

The paper is structured as follows: in section 2 related works are presented. We then provide a brief overview of modern automotive architecture, communication and functions (in sections 3 - 7), followed by a thorough investigation on the feasibility of carrying out vehicular live forensics (in sections 8 - 9). The paper finally concludes in section 10 with conclusions and final remarks.

2 Related Work

Most vehicular forensic procedures today mainly concentrate on crash/accident investigations and scene reconstruction. Traditionally, this used to be carried out by physically examining the vehicular modules, but since these are increasingly being transformed into electronic systems, digital examination is now required, too. Moreover, most modern vehicles are equipped with an *Event Data Recorder (EDR)* [2,3] module or colloquially *black box*. Data collected by the EDR units include pre-crash information such as pre-crash system state and acceleration, driver input, and post-crash warnings. This information is clearly suitable for accident investigation, but not for criminal ones as ongoing surveillance requires data other than the operational state of the vehicle and selective longer-term retention. Nilsson and Larson have investigated the feasibility of combining both physical and digital vehicular evidence [4], showing that such approach improves typical crime investigations. They also carried out a series of related studies, mainly concerned with the security of the in-vehicle networks and how to detect attacks against them [5]. However, the focus of our work is somewhat different in that we take a more active role in our forensic examination and try to observe, in real-time, the behaviours of drivers and passengers, taking advantage of the recently introduced advanced electronic components and functions in typical modern higher-end vehicles.

3 Intelligent Vehicles Technology

The term *Intelligent vehicle* generally comprises the ability of the vehicle to sense the surrounding environment and provide auxiliary information in which the driver or the vehicular control systems can make judgments and take suitable actions. These technologies mainly involve passenger safety, comfort and convenience. Most modern vehicles implementing telematics (e.g. navigation) and driver assistance functions (e.g. parking assist), can be considered intelligent in this sense. Evidently, these functions are very rapidly spreading while becoming common even in moderately priced vehicles. This has highly motivated this research since, for the best of our knowledge, no previous work has been undertaken to exclusively investigate these new sources of information vehicles can offer for digital forensics examiners. However, before discussing such applications and functions, we first briefly review basic general design and functional principles of automotive electronic systems.

4 Automotive Functional Domains

When electronic control systems were first used in the 1970's vehicles, individual functions were typically associated with separate ECU. Although this unified ECU-function association was feasible for basic vehicle operation (with minor economical implications), it quickly became apparent that networking the ECUs was required as the complexity of systems increased and information had to be exchanged among units. However, different parts of the vehicle have different requirements in terms of performance, transmission and bandwidth, and also have different regulatory and safety requirements. Vehicular electronic systems may hence be broadly divided into several functional domains [6]: *(1) Power train domain:* also called drivetrain, controls most engine functions, *(2) Chassis domain:* controls suspension, steering and braking, *(3) Body domain:* also called interior domain, controls basic comfort functions like the dashboard, lights, doors and windows; these applications are usually called multiplexed applications, *(4) Telematics & multimedia domain:* controls auxiliary functions such as GPS navigation, hands-free telephony, and video-based functions, *(5) Safety domain:* controls functions that improve passenger safety such as belt pretensioners and tyre pressure monitoring.

Communication in the power train, chassis and safety domains is required to be in real-time for obvious reasons (operation and safety), while communication in the telematics & multimedia needs to provide sufficiently high data rates capable of transmitting bulk multimedia data. Communication in the body domain, however, does not require high bandwidth and usually involves limited amounts of data. In this paper, we are interested in functions that can provide a forensically useful data about driver and passenger behaviour; such data is mostly generated by comfort and convenience functions within the telematics & multimedia domain, though some functions in the body and safety domains are also of interest, as will be discussed later.

5 Automotive Networks and Bus Systems

Early interconnection requirements between ECUs were initially addressed by point-to-point links. This approach, however, increased the inter-ECU links exponentially as the number of ECUs increased, which introduced many reliability, complexity, and economical implications. Consequently, automotive networks emerged to reduce the number of connections while improving overall reliability and efficiency. Generally, automotive networks are either event-triggered (where data is transmitted only when a particular event occurs) or time-triggered (where data is periodically transmitted in time slots) [7]. In an attempt to formalise the distinction between these networks, the society of Automotive Engineers (SAE) classified automotive networks into four main classes: *(1) Class A:* for functions requiring low data rate (up to 10kbps), such as lights, doors and windows. An example of class A is LIN network. *(2) Class B:* mostly for data exchange between ECUs and has data rate of up to 125 kbps. An example of class B is Low speed CAN network. *(3) Class C:* for functions demanding high data rate up to 1Mbps (most functions in the Power train and Chassis domains). An example of class C is High speed CAN network. *(4) Class D:* for functions requiring data rate of more than 1Mbps, such as most functions in the Telematics & Multimedia domain, and some functions in the safety domain. Example of class D are FlexRay and MOST networks.

We note that a typical vehicle today consists of a number of different interconnected networks, thus any information generated by any ECU can be received at any other ECU [8]. However, since ECUs are classified into functional domains and each domain may deploy a different network type, *gateways* are used for inter-domain communication. In the following subsections we provide a brief overview of an example network from each class; table 1 presents a summary comparison between these networks [9].

LIN. Local Interconnect Network (LIN) was founded in 1998 by the LIN Consortium [10] as an economical alternative for CAN bus system and is mainly targeted for non-critical functions in the body domain that usually exchange low-volume data and thus does not require high data rates; such data is also not required to be delivered in real-time. LIN is based on master-slave architecture and is a time-driven network. Using an unshielded copper single wire, LIN bus can extend up to 40m while connecting up to 16 nodes. Typical LIN applications include: rain sensor, sun roof, door locks and heating controls [11].

CAN. Controller Area Network (CAN) [12] is an event-driven automotive bus system developed by Bosch and released in 1986 (latest version is CAN 2.0 released in 1991). CAN is the most widely used automotive bus system, usually connecting ECUs of the body, power train and chassis domains, as well as inter-domain connections. There are two types of CAN: *(1) Low-speed CAN:* standardized in ISO11519-2 [13], supports data rate of up to 125kbit/s and mostly operates in the body domain for applications requiring slightly higher transmission rate than LIN; example applications include: mirror adjustment, seat

Table 1. Comparison between the most popular automotive networks

	LIN	Low-CAN	High-CAN	FlexRay	MOST
Class	Class A	Class B	Class C	Class C & D	Class D
Domain	Body	Body, Power Train, chassis	Power Train, chassis	Power train, Chassis, Telematics & Mult., Safety	Telematics and Multimedia
Standard	LIN Consortium	ISO 11519-2	ISO 1198	FlexRay Consortium	MOST Consortium
Max. Data rate	19.2 kbit/s	125 kbit/s	1 Mbit/s	20 Mbit/s	22.5 Mbit/s
Topology	Bus	Bus	Bus	Star (mostly)	Ring
Max. node no.	16	24	10	22 per bus/star	64
Applications	windows, doors	lights, wipers	Engine, Transmission	airbag	CD/DVD player
Control Mechanism	Time-driven	Event-driven	Event-driven	Time/Event-driven	Time/Event-driven

adjustment, and air-conditioning. *(2) High-speed CAN:* standardized in ISO11898 [14], supports data rate of up to 1 Mbit/s and mostly operates in the power train and chassis domains for applications requiring real-time transmission; example applications include: engine and transmission management.

FlexRay. Founded by the FlexRay consortium in 2000, FlexRay [15] was intended as an enhanced alternative to CAN. FlexRay was originally targeted for X-by-Wire systems which require higher transmission rates than what CAN typically supports. Unlike CAN, FlexRay is a time-triggered network (although event-triggering is supported) operating on TDMA (Time Division Multiple Access) basis, and is mainly used by applications in the power train and safety domains, while some applications in the body domain are also supported [9]. FlexRay is equipped with two transmission channels, each having a capacity of up to 10 Mbit/s and can transmit data in parallel, achieving an overall data rate of up to 20 Mbit/s. FlexRay supports point-to-point, bus, star and hybrid network topologies.

MOST. Recent years have witnessed a proliferation of in-vehicle multimedia-based applications. These applications usually require high bandwidth to support real-time delivery of the large multimedia data. As a result, the Media Oriented Systems Transport (MOST) bus system [16] was developed in 1998 and is today the most dominant automotive multimedia bus system. Unlike CAN (which only defines the physical and data link layers), MOST comprises all the OSI reference model layers and even provides various standard application interfaces for improved interoperability. MOST can connect up to 64 nodes in a ring topology with a maximum bandwidth of 22.5 Mbit/s using an optical bus (though recent

MOST revisions support even higher data rate). Data in MOST network is sent in 1,024 bits frames, which suits the demanding multimedia functions. MOST supports both time-driven and event-driven paradigms. Applications of MOST including audio-based (e.g. radio), video-based (e.g. DVD), and telematics.

6 Automotive Sensors

A typical vehicle integrates at least several hundred sensors (and actuators, although this is not a concern for the present paper), with increasing number of sensors even in economical vehicles to provide new safety, comfort and convenience functions. Typically, ECUs are built from microcontrollers which control actuators based on sensor inputs. In this paper, we are not concerned with technical sensor technology issues such as how sensor information is measured and the accuracy or reliability of measurements, but rather in either the raw sensor information or the output of the ECU microcontrollers based on information from those sensors; for a comprehensive discussion about automotive sensors, the reader is referred to, e.g., [17].

7 Advanced Automotive Applications

Currently, typical modern vehicles contain around 30–70 Electronic Control Units (ECU) [18], most of which are part of the power train and the chassis domains and thus usually connected by CAN buses. However, while different vehicles maintain approximately similar number of these essential ECUs, the number of ECUs in other domains (especially the telematics & multimedia and safety domains) significantly differ for different vehicle models and they are mostly what constitute the intelligent vehicle technology. In the following we discuss examples of functions integrated in most modern, intelligent vehicles. Most of these functions are connected via MOST or FlexRay networks with few exceptions for functions that may be implemented in the body domain (and hence are typically connected by LIN or CAN links).

Adaptive Cruise Control. One of the fundamental intelligent vehicle functions is Adaptive Cruise Control (ACC). Unlike static cruise, which fixes the traveling speed of the vehicle, in ACC, the vehicle senses its surrounding environment and adjusts its speed appropriately; advanced ACC systems can also access the navigation system, identify the current location and adhere to the speed limit of the corresponding roadway and respond to road conditions. ACC can be based on Radar (radio waves measurements), LADAR (laser measurements), or computer vision (image/video analysis) [19]. In Radar and LADAR based ACC, radio waves and laser beams, respectively, are emitted to measure the range (the distance between the hosting vehicle and the vehicle ahead) and the range rate (how fast the vehicle ahead is moving), and adapt the traveling speed accordingly. In vision-based ACC, a camera mounted behind the windshield or the front bumper captures video images of the front scene in which computer vision algorithms are applied on to estimate the range and range rate [20]. Note that there are a few variants of ACC, e.g. high-speed ACC, low-speed ACC etc. While all of these variants are based on the same basic principles as outlined above, some of them take more active roles, such as automatic steering.

Lane Keeping Assist. Lane Keeping Assist (LKA) is an application of Lane Departure Warning Systems (LDWS) and Road Departure Warning Systems (RDWS). Motivated by safety reasons, LKA is now a key function in intelligent vehicles. The most widely used approach of implementing LKA is by processing camera images for the road surface and identifying lane edges (usually represented by white dashed lines), then either warn the driver or automatically adjust the steering away from the lane edge; similar process is applied when departing from roads. Other approaches to implement LKA include roadway magnetic markers detection, and using digital GPS maps [19], but these are less commonly used since not all roadways are equipped with magnetic markers (which is extremely expensive), while GPS lane tracking does not always produce acceptably accurate measures and may also be based on inaccurate maps.

Parking Assist. Parking assist systems are rapidly becoming an expected feature. Implementations range from basic ultrasonic sensor alerts to an automated steering for parallel parking as introduced in Toyota's Intelligent Parking Assist (IPS) system in 2003. Usually, these systems have an integrated camera mounted at the rear bumper of the vehicle to provide a wide angle rear-view for the driver and can be accompanied with visual or audible manoeuvre instructions to guide the vehicle into parking spaces.

Blind Spot Monitoring. Between the driver's side view and the driver's rearview, there is an angle of restricted vision usually called the *blind spot*. For obvious safety reasons, when changing the lane, vehicles passing through the blind spot should be detected, which is accomplished by the Blind Spot Monitoring (BSM) systems. Such systems detect vehicles in the blind spot by Radar, LADAR or Ultrasonic emitters, with vision-based approaches (i.e. camera image processing) also becoming increasingly common. Most of these systems initiate warnings to the driver once a vehicle is detected in the blind spot, but future models may take a more active role to prevent collisions by automatically control the steering. Note that blind spot monitoring may also refer to systems that implement an adjustable side mirrors to reveal the blind spot to the driver, e.g. [21], but here we refer to the more advanced (and convenient) RF- and/or vision-based systems.

Head-up Display and Night Vision. Head-Up Display (HUD) technology was originally developed for aircrafts. HUD projects an image on a vehicle's front glass (in aviation applications this was originally a separate translucent pane), which will appear for the driver to be at the tip of the bonnet, and can be used to display the various information such as dashboard information or even navigation instructions. Beginning in the mid-1990s, General Motors (GM) used

HUD technology to enhance visibility at night by adding night vision functions to the HUD. In this technology, the front bumper of the vehicle is equipped with an infrared camera which provides enhanced night vision images of the road ahead and projects it for the driver. Infrared cameras detect objects by measuring the heat emitted from other vehicles, humans or animals. Recent trends use Near-Infrared (NIR) cameras instead which are also able to detect cold objects like trees and road signs [19]. However, the range of NIR is shorter and extends for only around 100m compared to around 500m in the case of the conventional (thermal) infrared cameras.

Telematics and Multimedia. Originally motivated by location-based services, telematics is now a more general term and comprises all wireless communication to and from the vehicle to exchange various types of information, including navigation, traffic warnings, vehicle-to-vehicle communication and, recently, mobile Internet and mobile TV. Telematics services have seamlessly found their way to intelligent vehicles becoming totally indispensable from them. However, it is not clear whether multimedia-based services should be classified under telematics and indeed there is a fine line between the two; for brevity, and to prevent confusion, we here merge them under a single class and assume that they use similar bus technology (typically MOST or FlexRay). Multimedia-based services involve the transmission of large (and sometimes real-time) data, which require high data rates; examples of multimedia applications include hands-free phones, CD/DVD players, radio and voice recognition.

Navigation. Automotive navigation systems are among the most essential telematics applications in modern vehicles and can be either integrated or standalone. Off-the-shelf (aftermarket) standalone navigation systems operate independently from other in-vehicle automotive components; this type of portable systems is largely irrelevant for our discussion since it can be easily removed or tampered with, although some integration with other components via, e.g., Bluetooth may occur. Built-in navigation systems, on the other hand, are often tightly integrated with other in-vehicle ECUs. In this case, navigation is not solely dependant on GPS technology, instead it takes advantage of its in-vehicle integration by receiving inputs from other automotive sensors; this is especially advantageous as GPS signals are not always available. Built-in navigation systems use the *Vehicle Speed Sensor (VSS)* or *tachometer sensor* to calculate the vehicle's speed, the *yaw rate sensor* to detect changes in direction, and GPS to determine the absolute direction movement of the vehicle. Integration also provides further benefits in applications such as *Adaptive Light Control*, automatically adjusting headlight settings to, e.g., anticipate turns, or simply by highlighting points of interest such as petrol stations in low-fuel situations.

Occupant Sensors. For safety reasons, it is important to detect the presence of occupants inside the vehicle. This is usually accomplished by mounting sensors under the seats to detect occupancy by measuring the pressure of an occupant's weight against the seat [22]. More advanced systems can even estimate the size of the occupant and consequently adjust the inflation force of the airbag in case of an accident since inflating the airbag with sufficiently high pressure can sometimes lead to severe injuries or even fatalities for children. Occupant detection can also be used for heating and seat belt alerts. However, rear seats may not always be equipped with such sensors, so another type of occupancy sensing primarily intended for security based on motion detectors is usually used [23]. These sensors can be based on infrared, ultrasonic, microwave, or radar and will detect any movements within the interior of the entire vehicle.

8 Live Forensics

Digital forensic examinations have rapidly become a routine procedure of crime and crime scene investigations even where the alleged criminal acts were not themselves technology-based. Although vehicular forensic procedures are slightly less mature than conventional digital forensics in personal computers and mobile (smart) phones, for example, we argue that the rich set of sensors and information obtainable from vehicles, as outlined above, can provide important evidence. Forensic examiners, therefore, are now starting to realise the importance of vehicular-based forensics and evidence. Moreover, as the same techniques can also be used, e.g., in (industrial) espionage, awareness of forensic techniques and counter-forensics in this domain are also becoming relevant. Typical forensic examinations are carried out either *offline* or *online* (live). Offline forensics involves examining the vehicle after an event while online forensics observe and report on the behaviour of a target in real-time. Note that this taxonomy may not agree with the literature where sometimes both offline and online forensics are assumed to take place post hoc and differ only by whether the vehicle is turned on or off, respectively, at the time of examination. Live forensics in this context is slightly different from surveillance as the latter may not always refer to exclusively observing criminals/suspects.

When adopting an online forensic approach, live data can be collected *actively* or *passively*. In either case, the system has to be *observed* appropriately before initiating the data collection process. In active live forensics, we have partial control over the system and can trigger functions to be executed without occupant knowledge. In passive live forensics, on the other hand, data are collected passively by intercepting traffic on vehicular networks. The observation process can be either hardware or software-based as discussed in sections 8.1 and 8.2, respectively. In both cases, data is collected by entities called *collectors*; while passive forensics may be approached by both software and hardware-based solutions, active forensics may only be feasible in a software-based approach owing to the (usually) limited time available to prepare a target vehicle for the hardware-based one.

As discussed in section 7, a typical intelligent vehicle integrates numerous functions usable for evidence collection and surveillance; this is a natural approach even for normal operation. For example, parking assist units are sometimes used by the automatic steel folding roof systems in convertibles to first monitor the area behind the vehicle and assesses whether folding the roof is possible. Similarly, we can observe and collect the output of relevant functions and draw conclusions about the behaviour of the occupants while using such data as evidence. We generally classify the functions we are interested in as vision-based and RF-based functions, noting that some functions can use a complementary vision-RF approach, or have different modes supporting either, while other functions based on neither vision or RF measurement can still provide useful information as shown in section 9:

(1) Vision-based functions: these are applications based on video streams (or still images) and employ computer vision algorithms – sometimes we are interested in the original video data rather than the processed results. Examples of these applications include: ACC, LKA, parking assist, blind spot monitoring, night vision, and some telematics applications. Vision-based applications are generally based on externally mounted cameras, which is especially useful to capture external criminal activities (e.g. exchanging/selling drugs), even allowing to capture evidence on associates of the target. Furthermore, newer telematics models may have built-in internal cameras (e.g. for video conferencing) that can capture a vehicle's interior.

(2) RF-based functions: similarly, these are applications based on wireless measurements such as ultrasonic, radar, LADAR, laser or Bluetooth. Unlike vision-based applications, here we are mostly interested in post-analysis of these measurements as raw RF measurements are typically not forensically meaningful.

8.1 Hardware-Based Live Forensics

The most straightforward solution for live forensics is to adopt a hardware-based data collection approach which involves installing special intercepting devices (*collectors*) around the vehicle to *observe* and collect the various types of data flowing through the vehicular networks. The collectors can be attached to ECU's or other components and capture outbound and/or inbound traffic. This information may then be locally stored inside the collectors or in a central location such as an entertainment system for later retrieval if sufficient local storage is available, or otherwise, the collectors can be configured to establish a private connection to an external location (i.e. federated network) for constant data transmission. This private network can, e.g., be setup through GSM/UMTS in cooperation with the carrier.

It is of utmost importance to carefully decide where to install these collectors, thus a good understanding of the data flow within the in-vehicle automotive system is required. Since different vehicle makes and even models have slightly different specifications, in this section we try to discuss the most attractive *observation loci* within the vehicle. As described above, vehicular systems contain several networks of different types that are interconnected by gateways, which can be considered the automotive equivalent of routers in conventional networks. Either a central gateway is used where all networks are connected to a single gateway (see figure 1(a)), or these networks are connected by several gateways (see figure 1(b)). In our live forensics examination, we are only interested in data generated by specific ECU's (mostly those that are part of MOST or FlexRay networks which correspond to functions in the body, telematics and safety domains), thus only those gateways connecting such networks need to be observed. However, in some cases, observing the gateways only may not be sufficient because in some applications we may also be interested in the raw ECU sensor readings (such as camera video/images) which may be inaccessible from gateways. For example, in vision-based blind spot monitoring application, the information relevant to the driver is whether there is an obstacle at the left/right side of the vehicle, but we are not interested in this information, we are only interested in the video/image that the corresponding sensors capture to be used to detect the presence of an obstacle (i.e. we are interested in the ECU input, while only the output is what normally sent through the gateway). Thus, in such cases, we may need to observe individual ECU's rather than gateways. Note, however, that observing gateways only may work for some applications where the input and the output are similar, such as parking assist where the parking camera transmits a live video stream to the driver.

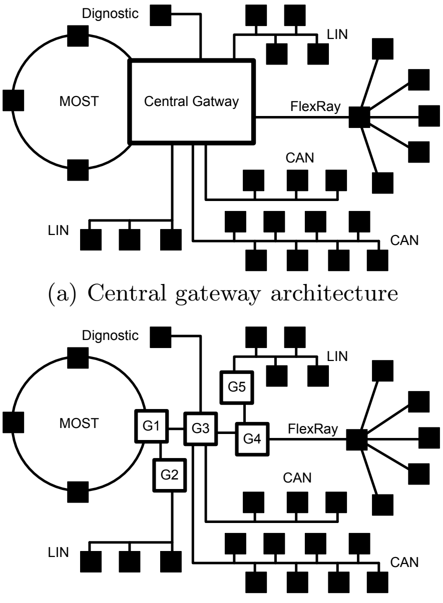

(a) Central gateway architecture

(b) Distributed gateway architecture

Fig. 1. Sample Automotive Network Architectures

8.2 Software-Based Live Forensics

Although by simply installing hardware collectors at particular ECU's or gateways, we will be able to collect live forensic data, such an approach may be limited due to aspects: *(1) Flexibility:* Since installation and removal of the hardware collectors need to be carried out manually and physically, they are inflexible in terms of reconfigurability and mobility; that is, once a devices is installed, it cannot be easily reconfigured or moved without physical intervention, which is not convenient or even (sometimes) possible. *(2) Installation:* The installation process of these devices will pose a serious challenge as locating and identifying the relevant ECU's or gateways is often difficult especially when some functions use information from several ECU's and sensors. Moreover, physical devices may be observable by an investigating target. *(3) Inspection:* the collectors will very likely collect large amount of possibly irrelevant data (such as channel management data); although this can be mitigated by using slightly more sophisticated collectors that filter observed traffic before interception, this introduces cost and efficiency implications.

Software based solutions, on the other hand, seem to alleviate these problems. Traditionally, the in-vehicle software (firmware) is updated manually via the vehicle's on-board diagnostic port. However, with the introduction of wireless communication, most manufacturers are now updating the firmware wirelessly, which, in turn, introduced several security concerns. Indeed, a recent work [24] showed that automotive networks are still lacking sufficient security measures. Thus, in our scenario, and following a particular set of legal procedures (see section 10), we can install the collectors as firmware updates with relative ease. These updates are then injected into the in-vehicle networks wirelessly and routed to the appropriate ECU.

Although software-based live forensics may be flexible and efficient, it poses a whole new class of compatibility and potentially safety issues. Unfortunately, most of the current software-based automotive solutions are proprietary and hardware dependant; thus, it may appear that unless we have knowledge of the automotive software and hardware architecture we are targeting, we will not be able to develop a software application to carry out our live forensics process, and even if we have such knowledge, we will be able to develop such software that will only work in the system it was developed for (lack of interoperability). However, these interoperability limitations (which also affect other automotive applications) have recently been realised and drove the leading automotive manufacturers and suppliers to establish an alliance for developing a standardized software architecture, named AUTOSAR.

AUTOSAR. AUTomotive Open System ARchitecture (AUTOSAR) is a newly established initiative by a number of leading automotive manufacturers and suppliers that jointly cooperated to develop a standardized automotive software architecture under the principle "cooperate on the standard, compete on the implementation". The first vehicle containing AUTOSAR components was launched in 2008 while a fully AUTOSAR supported vehicle is expected in 2010.

AUTOSAR aims to seamlessly separate applications from infrastructure so automotive applications developers do not have to be concerned about hardware peculiarities, which will greatly mitigate the complexity of integrating new and emerging automotive technologies. AUTOSAR covers all vehicle domains and functions from engine and transmission to wipers and lights. The main design principle of AUTOSAR is to abstract the automotive software development process and adopt a component-based model where applications are composed of software components that are all connected using a Virtual Functional Bus (VFB), which handles all communication requirements. AUTOSAR transforms ECUs into a layered architecture on top of the actual ECU hardware, as shown in figure 2 (a simplified view of the AUTOSAR layers). Below are brief descriptions of each layer:

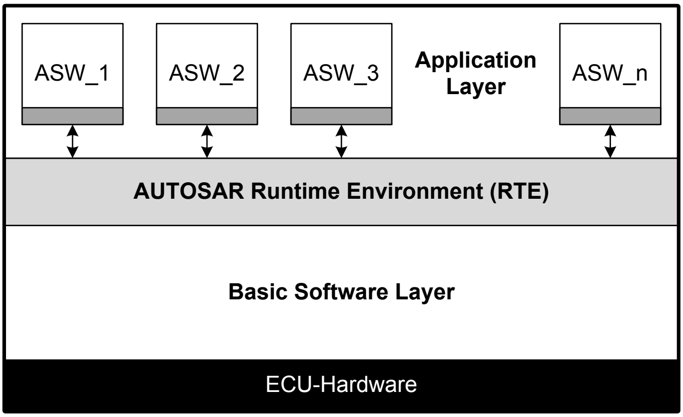

Fig. 2. AUTOSAR layered architecture

(1) AUTOSAR application layer: composed of a number of AUTOSAR software components (ASW). These components are not standardized (although their interfaces with the RTE are) and their implementation depends on the application functions. *(2) AUTOSAR Runtime Environment:* provides communication means to exchange information between the software components of the same ECU (intra-ECU) and with software components of other ECUs (inter-ECU). *(3) Basic software layer:* provides services to the AUTOSAR software components and contains both ECU independent (e.g. communication/network management) and ECU dependent (e.g. ECU abstraction) components. AUTOSAR standardises 63 basic software modules [25].

All software components are connected through the Virtual Functional Bus (VFB) which is implemented by the RTE at each ECU (VFB can be thought of as the concatenation of all RTE's). This paradigm potentially hides the underlying hardware from the application view, which, clearly, has advantageous consequences when collecting evidence for forensic examination. Thus an AUTOSAR-based collection tool will be compatible with all AUTOSAR supported vehicles. Furthermore, since the VFB allows seamless collection via different software components at different ECUs, a single live forensic software will be able to communicate with different software components and retrieve data from other applications and functions without having to be concerned with communication and other ECU-dependant issues.

Active Software-Based Live forensics. As discussed above, active live forensics appears feasible mainly when collectors are based on software, and is further facilitated by architectures such as AUTOSAR. An example of a typical application where active live forensics can be carried out is the vehicle's built-in hands-free telephony system. Although features and functions offered by the hands-free system may be different from a particular vehicle model to another, most recent models of hands-free system will synchronise some information with the phone it is paired with, including address books (contacts list) and call history. One benefit of this synchronisation process is allowing the driver to interact with the phone through the vehicle entertainment and communication system instead of the handset itself. This functionality is particularly useful for our live forensic investigation since it means that once the phone is paired with the hands-free system, the hands-free system can control it. Thus, an obvious active live forensic scenario is for the collector to initiate a phone call (without the knowledge of the driver) to a particular party (e.g. law enforcement) and carry out a live audio-based surveillance, the police can then cooperate with the carrier to suppress the relevant call charges. This can also occur in a side-band without affecting the ability to conduct further calls or in bursts.

We also note that the ability to scan for Bluetooth (or other RF such as 802.11) devices within a vehicle provides further potential to establish circumstantial evidence of the presence of individuals in a vehicle's proximity, even if, e.g., a passenger's mobile phone is never paired with the vehicle's communication system, allowing further tracking as reported in previous research [26].

9 Sensor Fusion

Forensic investigations can be significantly improved by fusing information from different sources (sensors). Many functions already implement sensor fusion as part of their normal operation, where two sensor measurements are fused, e.g. park assist uses ultrasonic and camera sensors. Similarly, while carrying out live forensic, we can fuse sensor data from even different functions that are not usually fused, such as video streams from blind spot monitoring with GPS measurements, where the location of the vehicle can be supported by visual images. Generally, however, data fusion is a post hoc process since it usually requires more resources than what the collectors are capable of. Below we discuss two applications of data fusion.

Visual Observation. Fusing video streams from different applications may result in a full view of the vehicle's surrounding environment. This is possible as the front view is captured by ACC, the side views by blind spot monitoring, and back view by parking assist cameras, while some vehicles provide further surround views. Note, however, that some of these cameras are only activated when the corresponding function is activated (e.g. the parking assist camera is only activated when the driver is trying to park); but obviously, active forensics can surmount this problem as it can actively control (activate/deactivate) the relevant functions.

Occupant Detection. As discussed in section 7, occupancy can be detected through existing sensors. However, further identifying the individuals on-board is even more desirable than just detecting their presence. While the approach of scanning for Bluetooth MAC addresses mentioned in section 8.2 may possibly identify the occupants passively, audio and, potentially, video recordings can provide further evidence even about individuals approaching or leaving the vehicle. Furthermore, In an active live forensic scenario, both the hands-free system and the occupant detection sensors can be associated such that if the occupant sensor detected a new occupant, the hands-free system automatically (and without the knowledge of the driver) initiates a pairing search to detect all MAC addresses in range. Note that hands-free search may detect Bluetooth devices of nearby vehicles or pedestrians and must hence be fused with occupant detection sensors information and repeated regularly, augmented by cameras where possible.

10 Discussion and Conclusion

The mechanisms (both active and passive) described in this paper have significant privacy and legal implications, yet while presenting this work we assume that such procedures are undertaken by law enforcement officials following appropriate procedures. We note that in some jurisdictions, it may not be necessary to obtain warrants, which is of particular relevance when persons other than the driver or vehicle owner are observed; this is, e.g., the case under the United Kingdom's Regulation of Investigatory Powers Act (2000).

In this paper, we presented a general overview of modern automotive systems and further discussed the various advanced functions resulting in what is commonly known today as an *Intelligent Vehicle*. We showed that functions available in modern automotive systems can significantly improve our live (real-time) digital forensic investigations. Most driver/passenger comfort and convenience functions such as telematics, parking assist and Adoptive Cruise Control (ACC) use multimedia sensors capturing the surrounding scene, which, if properly intercepted, can provide substantial evidence. Similarly, other sensors, like seat occupant sensors and hands-free phone systems, can be used for driver/passenger identification.

Future work will concentrate on characterising and fusing sensor data sources, while a natural extension to this work is to look at the feasibility of offline forensics (post hoc extraction of data) and investigate what kind of non-volatile data (other than Event Data Record (EDR) data, which is not always interesting or relevant for forensic investigations) that the vehicular system preserves and stores in-memory. Our expectation is that most of such data is not forensically relevant to investigating behavioural analysis of individuals in a court of law. However, we note that some functions may be capable of storing useful information as part of their normal operation, possibly with user interaction. For example, most navigation systems maintain historical records for previous destinations entered by the user in addition to a *favourite locations* list and a *home* location bookmark configured by the user; these records and configurations are likely to be non-volatile and can be easily retrieved at later times. Moreover, these systems may also contain information on *intended* movement, which is of particular interest if it can be communicated in real-time to investigators and enables anticipating target movements. Finally, future work will investigate counter-forensics mechanisms, which may also be relevant to investigate that vehicles such as hire cars have not been tampered with in anticipation of industrial espionage operations.

References

1. Wilwert, C., Navet, N., Song, Y., Simonot-Lion, F.: Design of Automotive X-by-Wire Systems. In: Zurawski, R. (ed.) The Industrial Communication Technology Handbook. CRC Press, Boca Raton (2005)
2. Singleton, N., Daily, J., Manes, G.: Automobile Event Data Recorder Forensics. In: Shenoi, S. (ed.) Advances in Digital Foreniscs IV. IFIP, vol. 285, pp. 261–272. Springer, Heidelberg (2008)
3. Daily, J., Singleton, N., Downing, B., Manes, G.: Light Vehicle Event Data Recorder Forensics. In: Advances in Computer and Information Sciences and Engineering, pp. 172–177 (2008)
4. Nilsson, D., Larson, U.: Combining Physical and Digital Evidence in Vehicle Environments. In: 3rd International Workshop on Systematic Approaches to Digital Forensic Engineering, pp. 10–14 (2008)
5. Nilsson, D., Larson, U.: Conducting Forensic Investigations of Cyber Attacks on Automobile in-Vehicle Network. In: e-Foreniscs 2008 (2008)
6. Navet, N., Simonot-Lion, F.: Review of Embedded Automotive Protocols. In: Automotive Embedded Systems Handbook. CRC Press, Boca Raton (2008)
7. Shaheen, S., Heffernan, D., Leen, G.: A Comparison of Emerging Time-Triggered Protocols for Automotive X-by-Wire Control Networks. Journal of Automobile Engineering 217(2), 12–22 (2002)
8. Leen, C., Heffernan, D., Dunne, A.: Digital Networks in the Automotive Vehicle. Computing and Control Journal 10(6), 257–266 (1999)
9. Dietsche, K.H. (ed.): Automotive Networking. Robert Bosch GmbH (2007)
10. LIN Consortium: LIN Specification Package, revision 2.1 (2006), `http://www.lin-subbus.org`
11. Schmid, M.: Automotive Bus Systems. Atmel Applications Journal 6, 29–32 (2006)
12. Robert Bosch GmbH: CAN Specification, Version 2.0 (1991)
13. International Standard Organization: Road Vehicles - Low Speed Serial Data Communication - Part 2: Low Speed Controller Area Network, ISO 11519-2 (1994)
14. International Standard Organization: Road Vehicles - Interchange of Digital Informaiton - Controller Aera Nework for High-speed Communication, ISO 11898 (1994)
15. FlexRay Consortium.: FlexRay Communications Systems, Protocol Specification, Version 2.1, Revision A. (2005), `www.flexray.com`
16. MOST Cooperation: MOST Specifications, revision 3.0 (2008), `http://www.mostnet.de`
17. Dietsche, K.H. (ed.): Automotive Sensors. Robert Bosch GmbH (2007)
18. Prosser, S.: Automotive Sensors: Past, Present and Future. Journal of Physics: Conference Series 76 (2007)

19. Bishop, R.: Intelligent Vehicle Technology and Trends. Artech House, Boston (2005)
20. Stein, G., Mano, O., Shashua, A.: Vision-based ACC with a Single Camera: Bounds on Range and Range Rate Accuracy. In: IEEE Intelligent Vehicle Symosium (2003)
21. Suggs, T.: Vehicle Blind Spot Monitoring System (Patent no. 6880941) (2005)
22. Henze, K., Baur, R.: Seat Occupancy Sensor (Patent no. 7595735) (2009)
23. Redfern, S.: A Radar Based Mass Movement Sensor for Automotive Security Applications. IEE Colloquium on Vehicle Security Systems, 5/1–5/3 (1993)
24. Nilsson, D., Larson, U.: Simulated Attacks on CAN Busses: Vehicle Virus. In: AsiaCSN 2008 (2008)
25. Voget, S., Golm, M., Sanchez, B., Stappert, F.: Application of the AUTOSAR Standard. In: Navet, N., Simonot-Lion, F. (eds.) Automotive Embedded Systems Handbook. CRC Press, Boca Raton (2008)
26. Al-Kuwari, S., Wolthusen, S.: Algorithms for Advanced Clandestine Tracking in Short-Range Ad Hoc Networks. In: MobiSec 2010. ICST. Springer, Heidelberg (2010)