Towards an Autonomous Vehicle Enabled Society: Cyber Attacks and Countermeasures[1]

Authors
Michael Haddrell, MSc student (Royal Holloway, 2016)
Keith M. Martin, ISG, Royal Holloway

Abstract
There are numerous exciting benefits of autonomous vehicles, which are very much at the frontier of development of vehicular technology. The level of cyber security provided by autonomous vehicles is, however, less well understood. Cyberattacks against autonomous vehicles could have grave implications for safety, privacy, reputation of manufacturers, and public perception. We examine the threats and vulnerabilities associated with autonomous vehicles, as well as discussing appropriate countermeasures. We argue that the successful embracing of autonomous vehicles in future transport systems will be best achieved by taking a sensible risk-management approach to tackling the potential cyber security threat.

Autonomous Vehicles
Over the years vehicles have become much more intelligent in a number of different ways. The latest wave of technological evolution is bringing vehicles that offer unprecedented levels of autonomy. Autonomous vehicles are capable of sensing their environment using sensors such as radar, laser-based lidar, GPS, cameras, compasses and wheel odometers. They can then take this information and make navigational decisions without human intervention.

The UK Parliamentary Office of Science and Technology has defined differing levels of autonomy (see Table 1). These range from vehicles with automated subsystems through to fully autonomous vehicles that are capable of completing an entire journey without any human input.

Fully autonomous vehicles have already been developed and are available, but thus far are only suitable for pedestrianised zones. Highly autonomous prototype vehicles are already taking part in road trials. It is estimated that it will be approximately 20 years before widely deployed fully autonomous vehicles and the necessary supporting infrastructure are a reality on public roads.

---

[1] This article is to be published online by Computer Weekly as part of the 2016 Royal Holloway information security thesis series. It is based on an MSc dissertation written as part of the MSc in Information Security at the ISG, Royal Holloway, University of London. The full MSc thesis is published on the ISG's website.

		Levels of autonomy	Existing examples
	Driver only	The vehicle is entirely under human control but may have some automated systems.	Cruise control, electronic stability control, anti-lock brakes.
	Driver assistance	The steering and/or acceleration are automated but the driver must control other functions.	Adaptive cruise control: distance to car in front is maintained. Parking assistant: Steering is automated, driver controls accelerator and brake.
	Partial autonomy	The driver does not control steering or acceleration but is expected to be attentive at all times and take back control instantaneously when required.	Adaptive cruise control with lane keeping. Traffic jam assistance.
	High autonomy	Vehicles are able to operate autonomously for some portions of the journey. Transfer of control back to the human driver happens with some warning.	Prototype vehicles.
	Full autonomy	The vehicle is capable of driving unaided for the entire journey with no human intervention – potentially without a human in the car.	None yet available for use on public roads

Table 1 - Levels of Vehicle Autonomy

The Benefits

The many potential benefits of autonomous vehicles include improvement in road safety, increased fuel efficiency and reduced congestion. They will also allow travellers to become more productive when on the move, as well as offering greater mobility to a wider range of individuals (for example, those with disabilities that make it hard to manually control a vehicle).

The Government wants the UK to become a world leader in autonomous vehicle technology. The Department for Transport has created an action plan in order to ensure that potential barriers to the introduction of autonomous vehicles do not prevent them from being tested and introduced. The UK Government has also provided £19 million to launch four autonomous vehicle schemes in different locations.

Existing Threats and Vulnerabilities

Autonomous vehicles are sophisticated systems containing computers that are connected to the Internet and one another. It would thus seem inevitable that, just as for other similar systems, attempts will be made by to infiltrate and compromise them.

Consideration of threats and vulnerabilities associated with fully autonomous vehicles also requires examination of the cyber security of contemporary 'driver only' and 'driver assistance' vehicles that are already prevalent on UK roads since these vehicles will provide the foundation for autonomous vehicles.

Modern vehicles currently contain more than one hundred electronic control units (ECUs) to assist with functions such as braking, speed (cruise) control and steering (parking assistance). These ECUs are networked and interact using an internal vehicle communication system such as the Controller Area Network (CAN) Bus, or other communication systems such as the Local Interconnect Network (LIN) Bus. This complex and growing array of interconnected ECUs is what facilitates the development of fully autonomous vehicles. Raw and interpreted sensor data provided by a central computer provides input to the ECUs via the internal vehicle communication system. This enables autonomous vehicles to make decisions on what actions to take whilst navigating to the chosen destination. Functionality and safety in both modern vehicles and future fully autonomous vehicles therefore depends on the integrity and availability of communications between these ECUs and the central computer.

As an example of the problems that can arise, in 2010 researchers highlighted the limitation of the CAN Bus protocol and exploited the fact that CAN protocol packets were not encrypted or authenticated when received by the ECUs. By reverse engineering the ECUs, listening to the network (sniffing) to perform protocol analysis, and by sending random and partially random packets to ECUs and observing the results (fuzzing), the researchers found a number of cyber security issues with the ECUs, the CAN Bus and CAN protocol packets that allowed them to control instrument display panels, locks and the braking system. They did, however, find that the car they examined contained two physically segmented networks – a high speed Bus (CANH) and a low speed Bus (CANL). The CAN Bus standard implicitly defines that the high-speed network is more trusted than the low speed network because the high-speed network connects the real-time safety critical components such as the ECUs for the engine and brakes. The low speed network contains components such as the radio and air-conditioning. According to the CAN Bus standard, gateways between the two networks can only be re-programmed from the high-speed network, thus preventing low-speed devices from attacking high-speed devices. However the researchers found ECUs connected to both networks and, whilst not a gateway, they re-programmed a dual connected ECU by uploading code to it from the low-speed network. This allowed them to send packets from the low speed network to the high-speed network, circumventing the network controls.

Attacks on vehicles of the type just described are most powerful if they can be conducted remotely, initiated from devices that contain remote access technologies such as a SIM card. One example of this is the CAN Hacking Tool (CHT), which costs £12 to create and allows an attacker to remotely inject packets into the CAN Bus. The CHT requires initial physical access to the vehicle for installation. However, modern vehicles are increasingly connected to the Internet through electronic systems designed for navigation, diagnostics, entertainment and safety monitoring. These systems provide potential external access points that could allow malicious actors remote access to the vehicle's internal communication system. Hackers have demonstrated that this type of exploitation is feasible by means of a broad range of attack vectors. These include a specially crafted music file, a Bluetooth stack vulnerability, a modem used for remote diagnostics, and even a bespoke digital audio broadcasting (DAB) radio station. Each attack allowed long distance full vehicle control, location tracking, in-cabin audio exfiltration, and theft of the vehicle.

New Threats and Vulnerabilities

Fully autonomous vehicles look set to inherit these existing cyber security issues. However, if designers are not careful, they are likely to introduce even more of their own.

One potential new area of vulnerability is sensors. Fully autonomous vehicles contain additional ECUs and rely on a centralised smart sensor infrastructure. One obvious attack vector is to manipulate a sensor in a way that could impact the physical behaviour of the vehicle by, for example, using reflective surfaces to affect distance perception of objects. However sensor attacks are unlikely to scale as well as attacks that manipulate the wider network that connects them.

Vehicular ad hoc networks (VANETs) have emerged to facilitate autonomous vehicles with their potential to improve traffic safety (for example collision avoidance), efficiency (congestion avoidance) and other added services. VANETs are realised by autonomous vehicles being able to communicate with their surrounding vehicles and/or roadside infrastructure (RSI). Figure 1 illustrates a VANET being used to relay accident information with the assistance of RSI that includes traffic lights.
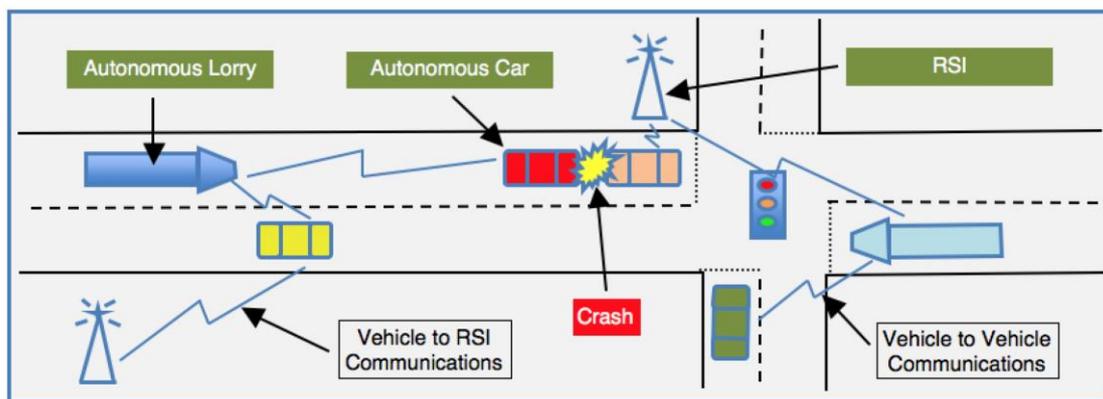

Figure 1 – Example of use of a VANET

A broad range of studies have considered the security problems and solutions related to VANET design. The attacks described by these studies tend to focus on the idea of a malicious or compromised node that affects the traffic flow (for example, causes congestion), safety (for example, causes an accident) or privacy (for example, tracks a vehicle) within the VANET by eavesdropping on legitimate communications and sending malicious messages.

Motivation for and Consequences of Cyberattack

While it is wise to consider the potential of cyberattack of autonomous vehicles and the underlying infrastructure necessary to support them, are laboratory-designed attacks likely to occur in real life?

The threat of such attacks naturally requires a motive and is hard to assess without specific risk analysis, which we shortly discuss. It is clear, however, that, as with cyberattacks on other technologies, this could be for reasons such as financial gain (theft), political protest (terrorism), surveillance or simply vandalism. It is also

possible that national governments and their associated cyber warfare programs are potential threat actors. A successful remote attacker making use of this increased attack surface could do more than simply apply the brakes; they could take full control and potentially direct vehicles to destinations of their choosing. If hackers targeted autonomous vehicles and their supporting infrastructure, there could be mass chaos.

A Risk Management Approach

Over the coming years we are going to see vehicles evolve from driver assistance vehicles towards fully autonomous vehicles. If this evolution mirrors other technologies then we are likely to see intense competition between manufacturers and pressure to be first to innovate. This in turn increases the potential danger of corner cutting and implementation mistakes, particularly with regard to cyber security, since past precedent suggest this is always lower down the development agenda than perhaps it should be.

However, threats, vulnerabilities and their associated risks, should not prevent innovative new ventures. There is no reason why the evolution of fully autonomous vehicles cannot benefit from our experience of risk assessment of existing technologies. When identified in advance, risks can be managed. Risk management in cyber security is an ongoing process that reduces risk by defining, analysing and controlling threats to, and vulnerabilities of, assets. It then identifies countermeasures to ensure that the associated risks and their potential impacts following an attack are mitigated to an acceptable level.

Asset: Anything that has value to the system or organisation.

Threat: A potential cause of an incident that may result in harm to the system or organisation.

Vulnerability: A weakness of an asset or group of assets that can be exploited by one or more threats.

Risk: The potential that a given threat will exploit vulnerabilities of an asset or group of assets and thereby cause harm to the system or organisation. Measured in terms of a combination of the likelihood of an event and its consequence.

Attack: A realised threat.

The detailed project report, whose findings this article is based on, follows a risk management process to examine the risks of cyberattacks against two autonomous vehicle applications proposed by the Traffic Research Laboratory (TRL). For each application, the risks were identified (labelled R1- R33) and countermeasures applied to successfully mitigate these risks, as summarised in Figure 2.
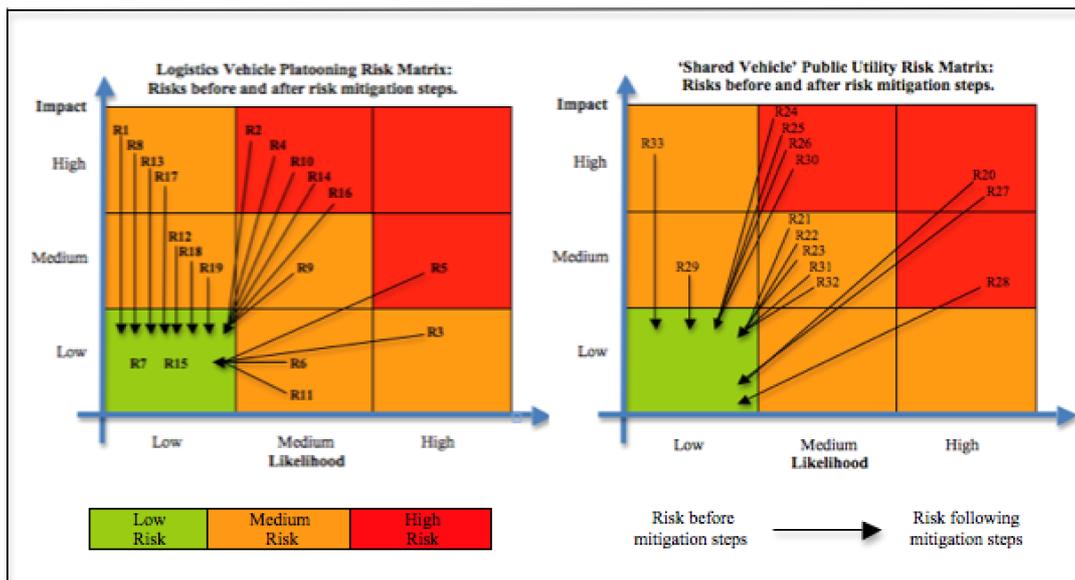
Whilst different autonomous vehicle applications share many risks, Figure 2 highlights that different applications can bring about new risks. In addition, some applications simply affect the risk level by, for example, changing the likelihood of a threat scenario occurring for that application.

The risk identified as R2 is an example of a shared risk. This risk relates to a remote attack via a vehicle's remote diagnostic system and is equally applicable to both applications. Similarly, risks introduced by vehicles having to rely on VANETs (R10 – R14) are also shared by both applications.

We now give an example of a risk that is applicable to both applications but is more of a risk against the 'Shared Vehicle' public utility application. This risk involves physically attaching a device onto the internal vehicle communication system in order to remotely control the vehicle. In the 'Logistics Vehicle Platooning' application this risk (R1) is not very likely to materialise because the attack requires internal access to the vehicle. However, in the 'Shared Vehicle' public utility application vehicles are openly accessible to paying customers and thus it is much easier for an attacker to place a device onto the internal vehicle communication system. This risk (R20) is therefore higher and appropriate countermeasures need to be developed.

Risks that are not applicable to both applications (for example risks R21-R26) are mainly due to the introduction of a vehicle occupant and the associated systems they require in the 'Shared Vehicle' public utility application. For example, occupants require entertainment systems and a system to request a vehicle to pick them up. These additions increase the remote surface attack area and so present new risks that are not applicable to the Logistics Vehicle Platooning application.

Recommendations
Whilst different applications of autonomous vehicles bring about different risks, the analysis that was conducted has made it possible to formulate a number of

recommendations to ensure autonomous vehicle development progresses. These will hopefully assist autonomous vehicle developers to enhance the cyber security of autonomous vehicles regardless of the application model being used.

It is recommended that:

- Further research should be undertaken to agree on cryptographic methods to prevent unauthorised commands from being sent to ECUs on the internal vehicle communication system.
- The low-speed internal vehicle communication system should be physically separated from the high-speed internal vehicle communication system containing the safety critical ECUs. Devices bridging these communication systems should be hardened and contain firewalls to allow only valid commands to traverse them.
- Remote access systems such as the diagnostic unit should employ cryptographically secure entity authentication techniques.
- Secure development practices such as the Security Development Lifecycle should be used when developing autonomous vehicle systems. There is no point in cryptographically securing the CAN Bus protocol if an attacker can exploit an ECU and steal the cryptographic key in order to send seemingly legitimate commands.
- VANETs should be designed carefully to provide resilience against jamming techniques and to provide the appropriate use of cryptography to prevent, for example, node tracking, node impersonation or unauthorised / modified messages. Mechanisms should enable malicious nodes to be revoked from the network.
- VANET standards should be agreed internationally without delay because these networks will provide the foundation for different models of autonomous vehicles and autonomous vehicle applications, bringing them together to operate seamlessly. This will result in autonomous vehicle applications being able to progress rapidly in a standardised manner, reducing complexity and costs.
- Audits (for example security penetration tests) by approved providers should test all autonomous vehicles, their systems and their supporting infrastructure designs, and implementations. This will ensure that vulnerabilities are detected before the system is activated. Designs that are secure in theory can be insecure in practice due to particular implementations and other reasons.
- Despite the low risks, reactive strategies should exist that assume the inevitability of attack against a vehicle and / or supporting infrastructure and so aim to detect, stop and allow further investigations into the attack methodology.
- Technology within the vehicle should also fail safely as per ISO/IEC 26262 and then services should (if possible) be restored through business continuity management and disaster recovery procedures.
- The Centre for the Protection of National Infrastructure should be engaged by organisations introducing autonomous vehicle applications to provide rapid and relevant cyber threat intelligence, and to create a coordination and incident handling function across autonomous vehicle and supporting infrastructure providers.

Concluding Remarks

Autonomous vehicles will be hugely beneficial to society but the implications of poor cyber security are grave. We need to get it right. Autonomous vehicles will require an innovative approach and cyber security concerns need not stifle innovation. A risk management approach enables an effective and efficient cyber security management framework for autonomous vehicles, supporting the objectives of introducing them by managing the risks. The approach is flexible and understands the context of each autonomous vehicle application, associated risk appetites and strategic direction of the organisations involved to identify proportionate controls in line with the operational desires and constraints. This will engender customer trust in the autonomous vehicle applications by ensuring that there will be a consistent, safe and quality service.

With the cyber security risks addressed, secure fully autonomous vehicles could reduce the current one million global traffic accident related deaths each year to, well, maybe even zero?

Biographies

Michael Haddrell is a current student at Royal Holloway, University of London whose interests include vulnerability research, penetration testing and cryptography. He is currently exploring options to work on a PhD in cyber security.

Keith Martin is a professor in, and former Director of, the Information Security Group at Royal Holloway, University of London. His current research interests include key management, cryptographic applications and securing lightweight networks. He is the author of the recently published "Everyday Cryptography" by Oxford University Press. As well as conventional teaching, Keith is a designer and module leader on Royal Holloway's distance learning MSc Information Security programme, and regularly presents to industrial audiences and schools.