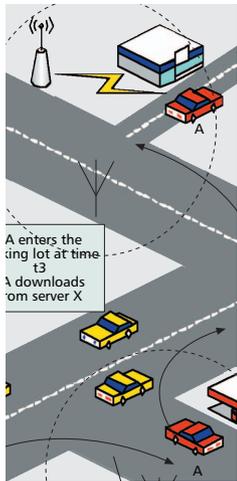


# SECURING VEHICULAR COMMUNICATIONS

MAXIM RAYA, PANOS PAPADIMITRATOS, AND JEAN-PIERRE HUBAUX, EPFL



The road to a successful introduction of vehicular communications has to pass through the analysis of potential security threats and the design of a robust security architecture able to cope with these threats. In this article the authors undertake this challenge.

## ABSTRACT

The road to a successful introduction of vehicular communications has to pass through the analysis of potential security threats and the design of a robust security architecture able to cope with these threats. In this article we undertake this challenge. In addition to providing a survey of related academic and industrial efforts, we also outline several open problems.

## INTRODUCTION

Initiatives to create safer and more efficient driving conditions have recently begun to draw strong support. Vehicular communications (VC) will play a central role in this effort, enabling a variety of applications for *safety*, *traffic efficiency*, *driver assistance*, and *infotainment*. For example, warnings for environmental hazards (e.g., ice on the pavement) or abrupt vehicle kinetic changes (e.g., emergency braking), traffic and road conditions (e.g., congestion or construction sites), and tourist information downloads will be provided by these systems.

Vehicular networking protocols will allow nodes, that is, vehicles or roadside infrastructure units, to communicate with each other over single or multiple hops. In other words, nodes will act both as end points and routers, with vehicular networks emerging as the first commercial instantiation of the *mobile ad hoc networking technology*.

The self-organizing operation and the unique features of VC are a double-edged sword: a rich set of tools are offered to drivers and authorities, but a formidable set of abuses and attacks becomes possible. Hence, the security of vehicular networks is indispensable, because otherwise these systems could make antisocial and criminal behavior easier, in ways that would actually jeopardize the benefits of their deployment. What makes VC security hard to achieve is the tight coupling between applications, with rigid requirements, and the networking fabric, as well as the societal, legal, and economical considerations. Solutions to this problem involve industry, governments, and academia, and can have a broad impact.

In this article we are specifically concerned with the following problem: how to design and

build vehicular communication protocols and systems that leave as little space as possible for misbehavior and abuse and, at the same time, remain resilient to ongoing attacks. We present an analysis of the vulnerabilities of vehicular networks and the salient challenges in securing their operation. Then we propose our architectural view of how VC can be secured, along with a brief (due to space limitations) overview of novel certificate revocation protocols tailored to the VC environment. Finally, we survey related works and discuss a few open issues in this emerging area of research.

## VULNERABILITIES AND CHALLENGES

### VULNERABILITIES

Any wireless-enabled device that runs a rogue version of the vehicular communication protocol stack poses a threat. We denote such rogue devices deviating from the defined protocols as *adversaries* or *attackers*.

The adoption of a variant of the widely deployed IEEE 802.11 protocol<sup>1</sup> by the vehicle manufacturers makes the attacker's task easier. And even possession of credentials cannot ensure alone the correct operation of the nodes. The effects of differing types of attackers (internal or external, rational or malicious, independent or colluding, persistent or random) can clearly differ. Here, rather than analyzing specific protocols, we are after a general exploration of VC vulnerabilities.

**Jamming** — The jammer deliberately generates interfering transmissions that prevent communication within their reception range. As the network coverage area (e.g., along a highway) can be well-defined, at least locally, jamming is a low-effort exploit opportunity. As Fig. 1 illustrates, an attacker can relatively easily, without compromising cryptographic mechanisms and with limited transmission power, partition the vehicular network.

**Forgery** — The correctness and timely receipt of application data is a major vulnerability. Figure 2 illustrates the rapid “contamination” of large portions of the vehicular network coverage area with false information where a single attacker forges and transmits false hazard warnings (e.g., ice formation on the pavement), which are taken up by all vehicles in both traffic streams.

<sup>1</sup> <http://grouper.ieee.org/groups/sc32/dsrc/>

**In-transit Traffic Tampering** — Any node acting as a relay can disrupt communications of other nodes: it can *drop* or *corrupt* messages, or *meaningfully modify* messages. In this way, the reception of valuable or even critical traffic notifications or safety messages can be manipulated. Moreover, attackers can *replay* messages (e.g., to illegitimately obtain services such as traversing a toll check point). In fact, tampering with in-transit messages may be simpler and more powerful than forgery attacks.

**Impersonation** — Message fabrication, alteration, and replay can also be used towards impersonation. Arguably, the source of messages, identified at each layer of the stack, may be of secondary importance. Often, it is not the source but the content (e.g., hazard warning) and the attributes of the message (freshness, locality, relevance to the receiver) that count the most. However, an impersonator can be a threat: consider, for example, an attacker masquerading as an emergency vehicle to mislead other vehicles to slow down and yield; or an adversary impersonating roadside units, spoofing service advertisements or safety messages.

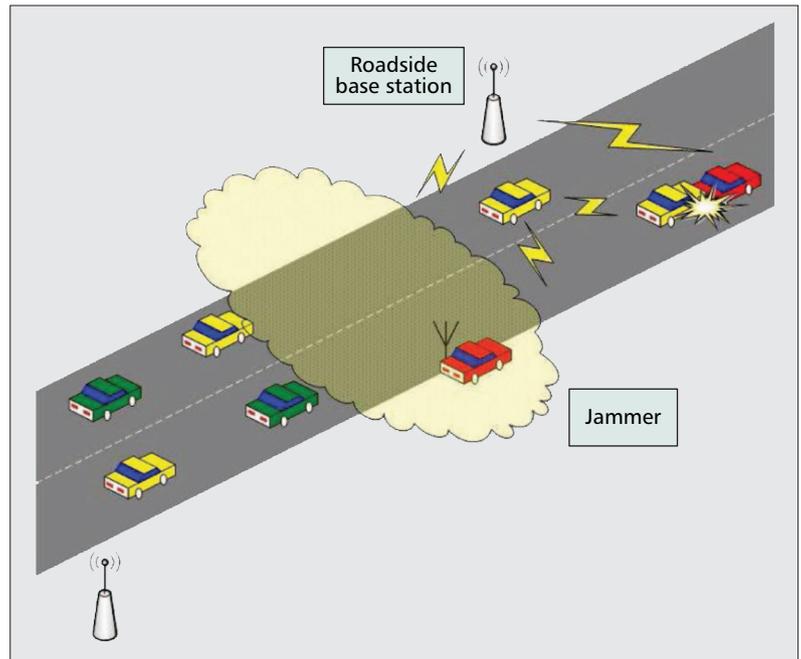
**Privacy Violation** — With vehicular networks deployed, the collection of vehicle-specific information from overheard vehicular communications will become particularly easy. Then inferences on the drivers' personal data could be made, and thus violate her or his *privacy*.<sup>2</sup> The vulnerability lies in the periodic and frequent vehicular network traffic: safety and traffic management messages, context-aware data access (e.g., maps, ferryboat schedules), transaction-based communications (e.g., automated payments, car diagnostics), or other control messages (e.g., over-the-air registration with local highway authorities). In all such occasions, messages will include, by default, information (e.g., time, location, vehicle identifier, technical description, trip details) that could precisely identify the originating node (vehicle) as well as the drivers' actions and preferences (Fig. 3).

**On-board Tampering** — Beyond abuse of the communication protocols, the attacker may select to tinker with data (e.g., velocity, location, status of vehicle parts) at their source, tampering with the on-board sensing and other hardware. In fact, it may be simpler to replace or by-pass the real-time clock or the wiring of a sensor, rather than modifying the binary code implementation of the data collection and communication protocols. Any VC security architecture should achieve a trade-off between robustness and cost due to tamper-proof hardware.

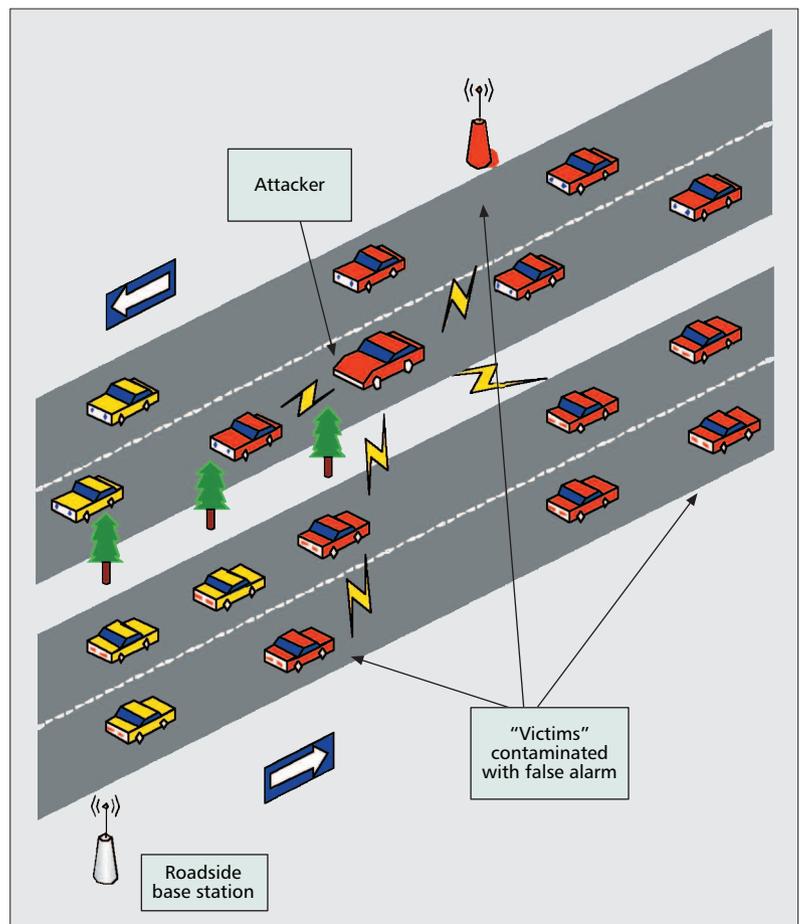
### CHALLENGES

The operational conditions, the constraints, and the user requirements for VC systems make security a challenging problem, with the most significant challenges specific to the VC discussed here.

<sup>2</sup> *Secrecy of personal data, such as those, for example, stored in repositories, and message confidentiality are not specific to VC only.*



■ Figure 1. Spectrum jamming.



■ Figure 2. Message forgery.

**Network Volatility** — The connectivity among nodes can often be highly transient and a one-time event. For example, two vehicles (nodes) traveling on a highway may remain within their

transceiver range, or within a few wireless hops, for a limited period of time. In other words, vehicular networks lack the relatively long-lived context and, possibly, the personal contact of the device users of a connection to a hot spot or the recurrent connection to an on-line service across the Internet. Hence password-based establishment of secure channels, gradual development of trust by enlarging a circle of trusted acquaintances, or secure communication only with a handful of endpoints may be impractical for securing VC.

**Liability vs. Privacy** — To make the problem more difficult, accountability and, eventually, liability of the vehicles and their drivers are required. Vehicular communication is envisioned as an excellent opportunity to obtain hard-to-refute data that can assist legal investigations (e.g., in the case of accidents). This implies that, to begin with, unambiguous identification of the vehicles as sources of messages should be possible. Moreover, context-specific information, such as coordinates, time intervals, and associated vehicles, should be possible to extract or reconstruct. But such requirements raise even stronger privacy concerns. This is even more so when drivers' biometrics are considered: biometrics, useful for enhancing vehicle access and control methods,

are highly private and unique data that cannot be reset or reassigned.

**Delay-Sensitive Applications** — Many of the envisioned safety and driver-assistance applications pose strict deadlines for message delivery or are time-sensitive. Security mechanisms must take these constraints into consideration and impose low processing and messaging overhead. Not only must protocols be lightweight, but also robust to clogging denial-of-service attacks. Otherwise, it would suffice for an adversary to generate a high volume of bogus messages and consume resources so that message delivery is delayed beyond the application requirements and thus, in practice, denied.

**Network Scale** — The scale of the network, with roughly a billion vehicles around the globe, is another challenge. This, combined with the multitude of authorities governing transportation systems, makes the design of a facility to provide cryptographic keys a challenge per se. A technically and perhaps politically convincing solution is a prerequisite for any security architecture.

**Heterogeneity** — The heterogeneity in VC technologies and the supported applications are additional challenges, especially taking into account the gradual deployment. With nodes possibly equipped with cellular transceivers, digital audio and Global Positioning System (GPS) or Galileo receivers, reliance on such external infrastructure should not be the weakest link in achieving security. For example, if GPS signaling can be spoofed, can the correctness of node coordinates and time accuracy be assumed? Second, with a range of applications with differing requirements, security solutions must retain *flexibility*, yet remain *efficient* and *interoperable*.

## SECURITY ARCHITECTURE

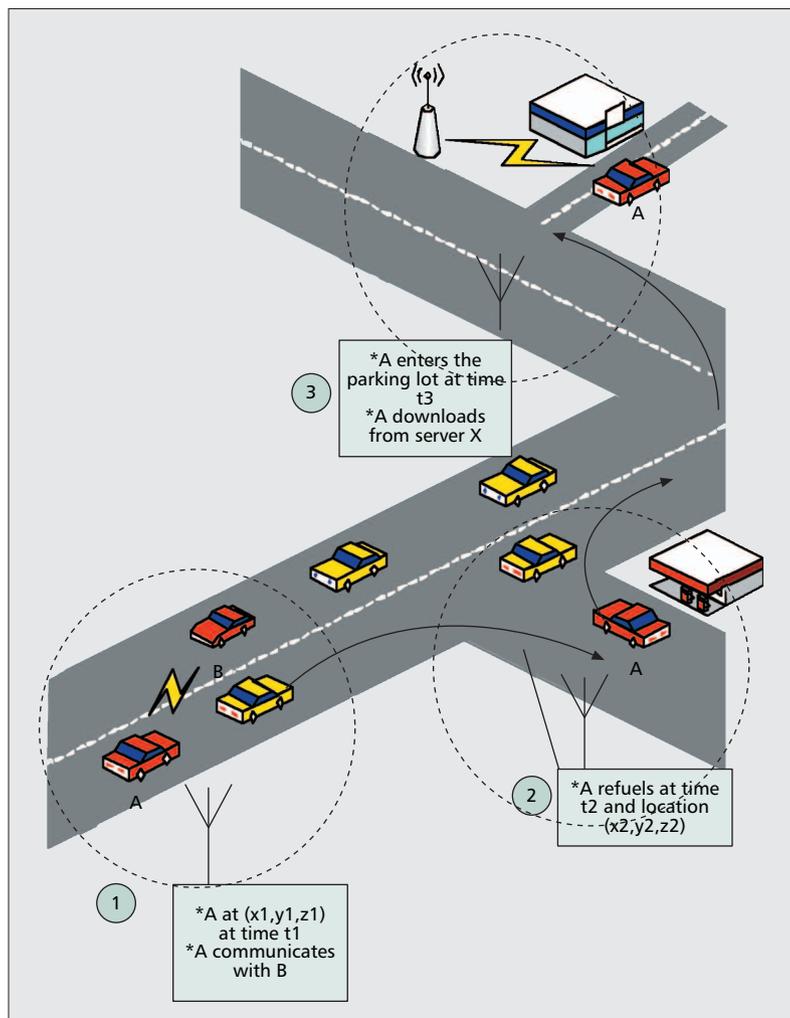
In this section we present the components needed to protect VC against a wide range of threats, some of which are described in the previous section. We also aim at providing an authentication, authorization, accounting (AAA) framework for VC. Figure 4 depicts the general architecture, the components of which are described next.

### SECURITY HARDWARE

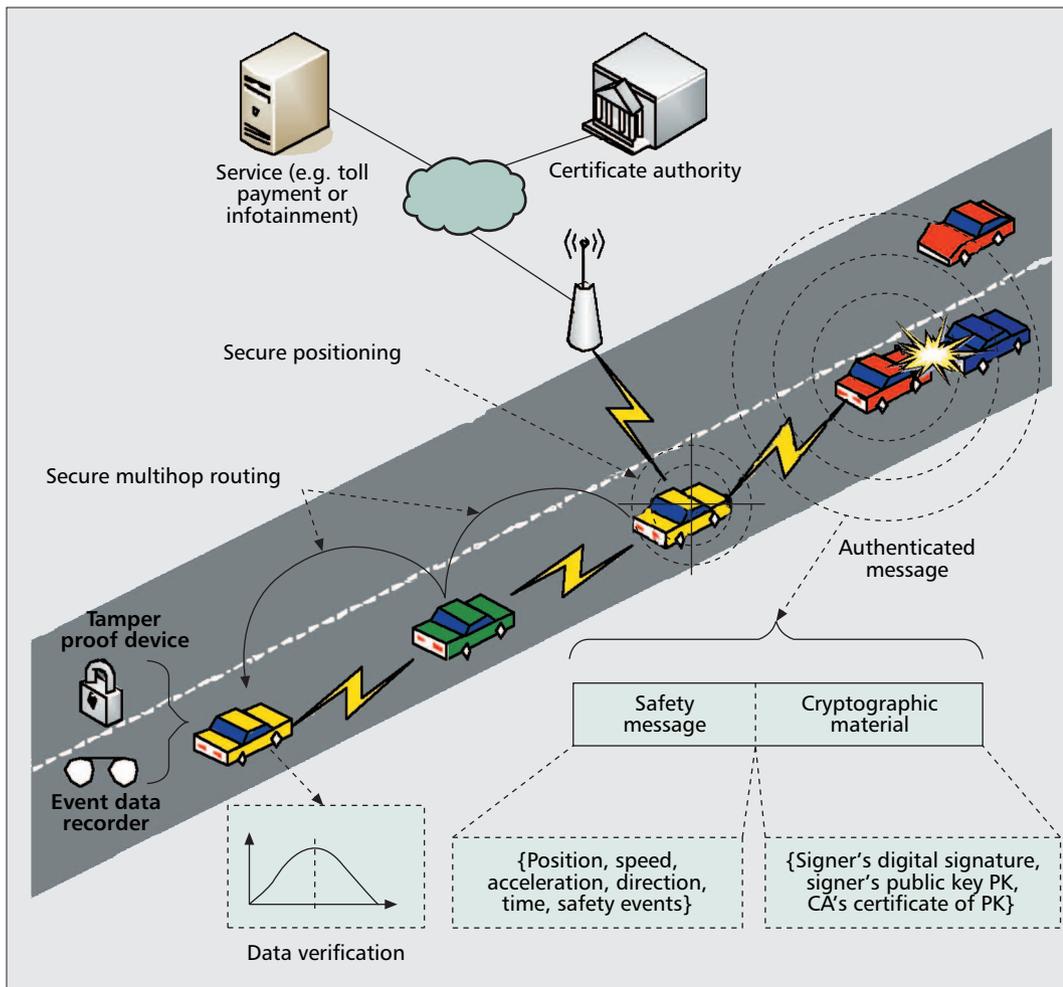
Among the vehicle onboard equipment, there should be two hardware modules needed for security, namely, the event data recorder (EDR) and the tamper-proof device (TPD). Whereas the EDR only provides tamper-proof storage, the TPD also possesses cryptographic processing capabilities.

The EDR will be responsible for recording the vehicle's critical data, such as position, speed, time, and so forth, during emergency events, and is similar to an airplane's black box. These data will help in accident reconstruction and the attribution of liability. EDRs are already installed in many road vehicles, especially trucks. These can be extended to also record the safety messages received during critical events.

The car electronics, especially the data bus system, are easily accessible by the owner or by



■ Figure 3. Vehicle tracking.



■ Figure 4. Overview of the security architecture.

a mechanic. Hence the cryptographic keys of a vehicle need proper hardware protection, namely, a TPD. The TPD will take care of storing all the cryptographic material and performing cryptographic operations, especially signing and verifying safety messages. By binding a set of cryptographic keys to a given vehicle, the TPD guarantees the accountability property as long as it remains inside the vehicle. The TPD has to be as independent as possible from its external environment; hence, it should include its own clock and have a battery that is periodically recharged from the vehicle's electric circuits. Yet, despite all these "features," the TPD will still suffer from the fact that it cannot control the correctness of the data it receives. This may result in the TPD signing messages with bogus data. The solution to this problem is briefly described in the "Authentication" subsection.

A major obstacle to the adoption of TPDs is their high cost. But current products are mainly intended for computation-hungry financial applications. Hence, there are several factors that can facilitate the introduction of TPDs in vehicles:

- The creation of a "lighter" version of TPDs
- The leverage on the building-up expertise for vehicular EDRs
- The economy of scale that will drive costs significantly lower

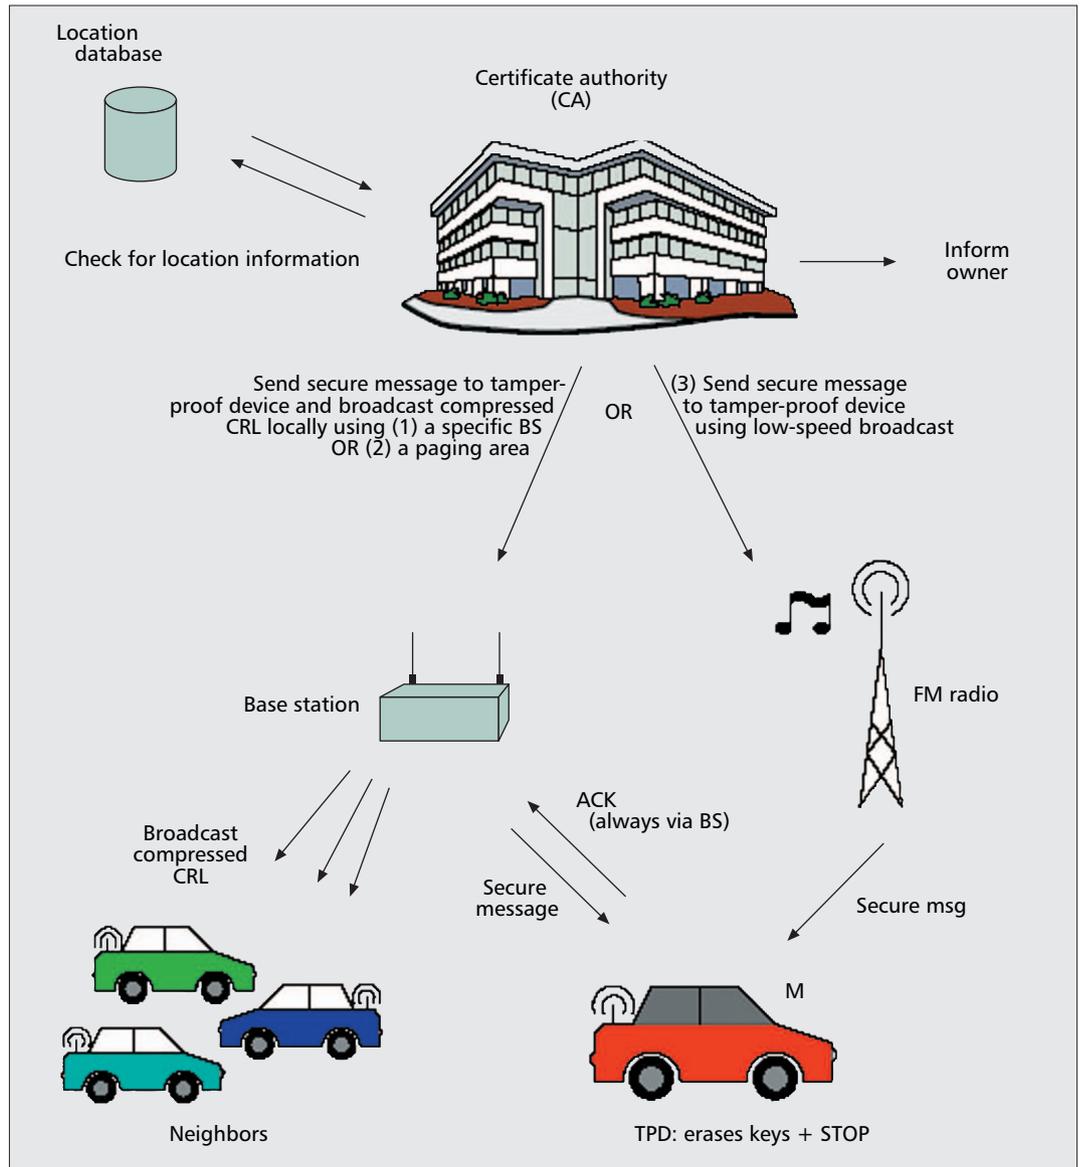
## VEHICULAR PUBLIC KEY INFRASTRUCTURE

The huge number of vehicles registered in different countries and traveling long distances, well beyond their registration regions, requires a robust and scalable key management scheme. The involvement of authorities in vehicle registration implies the need for a certain level of centralization. Communication via base stations (as in cellular networks) is not enough for VC, mainly because vehicles need to authenticate themselves not only to base stations, but also to each other (without invoking any server), which creates a problem of scalability. In addition, symmetric cryptography does not provide the nonrepudiation property that allows the accountability of drivers' actions (e.g., in the case of accident reconstruction or finding the originators of *forgery* attacks). Hence, the use of public key cryptography is a more suitable (if not the only) option for deploying VC security.

This implies the need for a Vehicular Public Key Infrastructure (VPKI) where Certificate Authorities (CAs) will issue certified public/private key pairs to vehicles (with many pairs per vehicle for privacy reasons, as is explained below). Similarly to current vehicle registration authorities, there will be several CAs, each corresponding to a given region (e.g., country, state, metropolitan area, etc.). Other candidates for

The huge number of vehicles registered in different countries and traveling long distances, well beyond their registration regions, requires a robust and scalable key management scheme. The involvement of authorities in vehicle registration implies the need for a certain level of centralization.

The fundamental security functions in VC will consist in authenticating the origin of a data packet. Authentication and the inherent integrity property counter the in-transit traffic tampering and impersonation vulnerabilities. In addition, authentication also helps to control the authorization levels of vehicles.



■ Figure 5. Revocation protocol of the tamper-proof device (RTPD).

taking the role of CAs are car manufacturers. In any of the two cases, the different CAs will have to be cross-certified so that vehicles from different regions or different manufacturers can authenticate each other. This will require each vehicle to store the public keys of all the CAs whose certificates it may need to verify. Alternately, in the case where CAs are regional authorities, vehicles may request new public/private key pairs delivered by the foreign region<sup>3</sup> they enter.

### AUTHENTICATION

The fundamental security functions in VC will consist in authenticating the origin of a data packet. Authentication and the inherent integrity property counter the *in-transit traffic tampering* and *impersonation* vulnerabilities. In addition, authentication also helps to control the authorization levels of vehicles.

To authenticate each other, vehicles will sign each message with their private key and attach the corresponding certificate. Thus, when another

vehicle receives this message, it verifies the key used to sign the message and, once this is done correctly, it verifies the message. To reduce the security overhead, the common approach is to use Elliptic Curve Cryptography (ECC) — the most compact public key cryptosystem thus far. But it is possible to reduce this overhead by signing only critical messages (e.g., with accident warnings) or one in every few messages (the frequency and redundancy of messages can allow this). In addition, given the frequency of safety message broadcasts (typically, every 300 ms), a vehicle can ignore redundant messages.

### CERTIFICATE REVOCATION

The advantages of using a PKI for VC are accompanied by some challenging problems, notably, certificate revocation. For example, the certificates of a detected attacker or malfunctioning device have to be revoked, that is, it should not be able to use its keys or, if it still does, vehicles verifying them should be made aware of their invalidity.

<sup>3</sup> In this context, "foreign" means a region different from a vehicle's home region.

Features	Network type			
	Cellular, WLAN	Sensor Networks	P2P (PGP)	VANET
Key Management	symmetric, centralized	symmetric, centralized	asymmetric, decentralized	asymmetric, multiple authorities
Authentication	authentication server	pairwise symmetric	digital signatures, web of trust	digital signatures, CA certificates
Revocation	directly by the operator	distributed voting	counter-certificates	short-lived certificates; CRLs
Privacy	temporary identifiers	NA	anonymizing services	preloaded keys
Positioning	triangulation with base stations	triangulation with beacons	NA	open problem

**Table 1.** Comparison of different network types with respect to security problems. It should be noted here that there exist several mechanisms proposed for some network types, but we consider the most widely adopted of these. Thus, for example, we took Pretty Good Privacy (PGP) as a representative example of Peer-to-Peer (P2P) security in the Internet.

The most common way to revoke certificates is the distribution of certificate revocation lists (CRLs) that contain the most recently revoked certificates; CRLs are provided when infrastructure is available. In addition, using short-lived certificates automatically revokes keys. These are the methods proposed in the IEEE P1609.2 standard [1]. But there are several drawbacks to this approach. First, CRLs can be very long due to the enormous number of vehicles and their high mobility (meaning that a vehicle can encounter a high number of vehicles when traveling, especially over long distances). Second, the short lifetime of certificates still creates a vulnerability window. Last but not least, the availability of an infrastructure will not be pervasive, especially in the first years of deployment.

To avoid the above shortcomings, we have designed a specific solution. It includes a set of revocation protocols, namely, Revocation Protocol of the Tamper-Proof Device (RTPD), Revocation protocol using Compressed Certificate Revocation Lists (RCCRL), and Distributed Revocation Protocol (DRP). In the following, we present the details of RTPD, as illustrated in Fig. 5, and only outline the main features of RCCRL and DRP (due to space constraints). In RTPD, once the CA has decided to revoke all the keys of a given vehicle  $M$ , it sends to it a revocation message encrypted with the vehicle's public key. After the message is received and decrypted by the TPD of the vehicle, the TPD erases all the keys and stops signing safety messages. Then it sends an ACK to the CA. All the communications between the CA and the vehicle take place in this case via base stations. In fact, the CA has to know the vehicle's location in order to select the base station through which it will send the revocation message. If it does not know the exact location, it retrieves the most recent locations of the vehicle from a location database and defines a paging area with base stations covering these locations. Then it multicasts the revocation message to all these base stations. In the case when there are no recent location entries or the ACK is not received after a timeout, the CA broadcasts the revocation

message, for example, via the low-speed FM radio on a nationwide scale or via satellite.

The RCCRL protocol is used when the CA wants to revoke only a subset of a vehicle's keys or when the TPD of the target vehicle is unreachable (e.g., by jamming or by tampering of the device). Given the expected large size of CRLs in vehicular ad hoc networks (VANETs), the key idea in RCCRL is to use Bloom filters — a probabilistic data structure used to test whether an element is a member of a set. Thus, the size of a CCRL will be only a few KB. RCCRL also relies on the availability of infrastructure that broadcasts the CCRLs once every 10 min. Compared to RTPD, RCCRL has the special feature of warning the neighbors of a revoked vehicle as they also receive the CCRLs.

The DRP protocol is used in the pure ad hoc mode whereby vehicles accumulate accusations against misbehaving vehicles, evaluate them using a reputation system, and, in case misbehavior is detected, report them to the CA once a connection is available. Unlike RTPD and RCCRL, the revocation in DRP is triggered by the neighbors of a vehicle upon the detection of misbehavior. Mechanisms for the detection of malicious data [2] can be leveraged to spot vehicles generating these data (since all messages are signed).

## PRIVACY

To address *privacy* vulnerability, we propose using a set of anonymous keys that change frequently (every couple of minutes) according to the driving speed. Each key expires after its usage; only one key can be used at a time. These keys are preloaded in the vehicle's TPD for a long duration, e.g., until the next yearly checkup; the TPD takes care of all the operations related to key management and usage. Each key is certified by the issuing CA and has a short lifetime (e.g., a specific week of the year). In addition, it can be tracked back to the real identity of the vehicle — the Electronic License Plate (ELP) — in case law enforcement necessitates this and only after obtaining a permission from a judge. This conditional anonymity will help determine the liability of drivers in the case of accidents.

There are many completed and ongoing projects on VC all over the world. Examples include the Berkeley PATH project in the United States and the German project Fleetnet. Yet none of these early projects has considered the security aspects of VC.

The downside of this approach is the necessity for storage space for all the keys for one year, but these can fit in only a few Mbytes [3].

In the case of infotainment applications in which vehicles communicate with the infrastructure, the CARAVAN scheme [4] allows vehicles to preserve their privacy by forming groups in which the group leader acts as a proxy on behalf of all group members that access the infrastructure. When the vehicles do not have to access the infrastructure, they remain silent, thus preventing eavesdroppers from tracking their pseudonyms.

## STATE OF THE ART ACADEMIC RESEARCH

The research on VC security is just beginning, with a few pioneer papers published thus far. In [5], Blum and Eskandarian describe a security architecture for VC intended mainly to counter the so-called “intelligent collisions” (meaning that they are intentionally caused). But this is only one type of attack, and building the security architecture requires awareness of as many potential threats as possible. They propose the use of a PKI and a virtual infrastructure where cluster-heads are responsible for reliably disseminating messages (by a sequential unicast instead of broadcast) after digitally signing them; this approach creates bottlenecks at cluster-heads in addition to high security overhead. Gerlach [6] describes the security concepts for vehicular networks. Hubaux et al. [7] take a different perspective of VC security and focus on privacy and secure positioning issues. They point out the importance of the trade-off between liability and anonymity and also introduce Electronic License Plates (ELPs), unique electronic identities for vehicles. Parno and Perrig [8] discuss the challenges, adversary types, and some attacks encountered in vehicular networks; they also describe several security mechanisms that can be useful in securing these networks. Raya and Hubaux [3] describe a full security and privacy framework for VANETs with primary simulation evaluations of the security overhead. El Zarki et al. [9] describe an infrastructure for VC and briefly mention some related security issues and possible solutions.

Table 1 summarizes the mechanisms used to provide security features in VC and compares them with other network types that are broadly addressed in the literature. We can see that the distinctive properties of VANETs, notably scale and high mobility, justify the need for, as well as the opportunity of, using novel solutions compared to other network types.

## INDUSTRIAL PROJECTS

There are many completed and ongoing projects on VC all over the world. Examples include the Berkeley PATH project in the United States and the German project Fleetnet. Yet none of these early projects has considered the security aspects of VC. To bridge this gap, new projects are allocating part of their resources to investigate security issues. In the following, we provide an overview of the most relevant ones.

The IEEE P1609.2 standard [1] is part of the DSRC standards for VC supported by the U.S. Vehicle Safety Communication Consortium (VSCC). It proposes using asymmetric cryptography to sign safety messages with frequently changing keys so that anonymity is preserved. There is no mechanism proposed for certificate revocation. Instead, certificates have short lifetimes and are periodically requested by vehicles through roadside base stations, implying the need for a pervasive infrastructure.

In Europe, VC security is partially considered within the projects Network on Wheels (NoW) and Global System for Telematics (GST) as well as by the Car2Car Communication Consortium (C2C-CC). It is being fully addressed by the new European project SEVECOM (that stands for Secure Vehicular Communications), which focuses on providing a full definition and implementation of security requirements for VC.

## OPEN PROBLEMS

In addition to the main building blocks presented earlier, there remains a set of unexplored problems directly related to VC security. In this section we outline the most important of these problems.

**Secure Positioning** — In VC, position is one of the most important data for vehicles. Each vehicle needs to know not only its own position but also those of other vehicles in its neighborhood. GPS signals are weak, can be spoofed, and are prone to jamming. Moreover, vehicles can intentionally lie about their positions. Hence, the need for a secure positioning system that will also support the accountability and authorization properties, frequently related to a vehicle’s position.

**Data Verification** — This helps to prevent the forging attacks illustrated in Fig. 2. This can be achieved by a *data correlation* mechanism that compares all collected data regarding a given event. A first example of such a mechanism is presented in [2], where the vehicle has a model to which it compares received data before classifying it as truthful, malicious, or unintentionally incorrect.

**DoS Resilience** — DoS attacks, and especially *jamming*, are relatively simple to mount, yet their effects can be devastating. Existing solutions such as frequency hopping do not completely solve the problem. The use of multiple radio transceivers, operating in disjoint frequency bands, can be a feasible approach.

## CONCLUSION

In this article we have described the problems that characterize the security of vehicular networks and sketched possible solutions. As we have seen, some of these solutions can leverage on existing security techniques. However, we also have stressed that vehicular communications exhibit unique security challenges, induced by the high speed and sporadic connectivity of the vehicles (especially with the infrastructure), the high relevance of their geographic location, the tension between liability and privacy, and the

huge scale and very gradual deployment of the network. Only a coordinated effort of all parties involved (vehicle manufacturers, transportation authorities, law enforcement agencies, insurance companies, and academic researchers) will make it possible to devise a solution that is compliant with the demanding requirements of this fascinating area.

Interested readers can visit <http://ivc.epfl.ch> for more information on this topic.

### ACKNOWLEDGMENTS

We would like to thank Virgil Gligor, Rainer Kroh, and Tim Leinmüller for their helpful feedback on earlier versions of this work.

### REFERENCES

- [1] "IEEE P1609.2 Version 1 — Standard for Wireless Access in Vehicular Environments: Security Services for Applications and Management Messages," in development, 2006.
- [2] P. Golle, D. Greene, and J. Staddon, "Detecting and Correcting Malicious Data in VANETs," *Wksp. Vehic. Ad hoc Networks (VANET)*, 2004.
- [3] M. Raya and J.-P. Hubaux, "The Security of Vehicular Ad Hoc Networks," *Wksp. Security in Ad hoc and Sensor Networks (SASN)*, 2005.
- [4] K. Sampigethaya *et al.*, "CARAVAN: Providing Location Privacy for VANET," *Wksp. Embedded Security in Cars (ESCAR)*, 2005.
- [5] J. Blum and A. Eskandarian, "The Threat of Intelligent Collisions," *IT Professional*, vol. 6, no. 1, Jan.–Feb. 2004, pp. 24–29.
- [6] M. Gerlach, "VaneSe: An Approach to VANET Security," *V2VCOM*, 2005.

- [7] J.-P. Hubaux, S. Capkun, and J. Luo, "The Security and Privacy of Smart Vehicles," *IEEE Security and Privacy Mag.*, vol. 2, no. 3, May–June 2004, pp. 49–55.
- [8] B. Parno and A. Perrig, "Challenges in Securing Vehicular Networks," *Wksp. Hot Topics in Networks (HotNets-IV)*, 2005.
- [9] M. El Zarki *et al.*, "Security Issues in a Future Vehicular Network," *Euro. Wireless*, 2002.

### BIOGRAPHIES

MAXIM RAYA ([maxim.raya@epfl.ch](mailto:maxim.raya@epfl.ch)) received a B.Eng. degree in computer and communications engineering in 2002 from the American University of Beirut, Lebanon. He is currently pursuing his Ph.D. studies at EPFL. His research interests are in the area of security in wireless networks, and especially vehicular networks. His web site is <http://people.epfl.ch/maxim.raya>.

PANAGIOTIS PAPADIMITRATOS ([panos.papadimitratos@epfl.ch](mailto:panos.papadimitratos@epfl.ch)) received his Ph.D. degree in electrical and computer engineering from Cornell University Ithaca, NY, in 2005. He joined the Department of Electrical and Computer Engineering at Virginia Tech, Blacksburg, VA, as a research associate. He is currently a senior researcher with the School of Computer and Communication Sciences at EPFL. His research is concerned with networking protocols, network security, ad hoc and sensor networks, and wireless and mobile systems. His web site is <http://people.epfl.ch/panos.papadimitratos>.

JEAN-PIERRE HUBAUX ([jean-pierre.hubaux@epfl.ch](mailto:jean-pierre.hubaux@epfl.ch)) is a full professor at EPFL. His research interests are mobile networking and computing, notably, the security and cooperation in fully self-organized mobile ad hoc networks. He is an Associate Editor of *IEEE Transactions on Mobile Computing and Foundations and Trends in Networking*. He has held visiting positions at the IBM T. J. Watson Research Center and at the University of California at Berkeley. His Web site is <http://people.epfl.ch/jean-pierre.hubaux>.