Secure Crash Reporting in Vehicular Ad hoc Networks

Sumair Ur Rahman and Urs Hengartner
David R. Cheriton School of Computer Science
University of Waterloo
Waterloo ON, N2L 3G1, Canada
{surrahman,uhengart}@cs.uwaterloo.ca

***Abstract*—We present AutoCore, an automated crash reporting application that uses VANETs (Vehicular Ad hoc NETworks) to provide authenticated digital video and telemetry data. This data is recorded by vehicles either involved in or at the scene of a crash and can be used by investigators to reconstruct the events that lead up to the crash. To secure this application, we present a security infrastructure that extends the state of the art in VANET security. In particular, the contributions of this infrastructure include (a) the concept of Road-worthiness Certificates, (b) use of these certificates in a practical scheme for the distribution of cryptographic vehicle credentials issued by regional transportation authorities, (c) a decentralized scheme for conditionally anonymous, inter-vehicle communication, (d) efficient support for the roaming of vehicles between different transportation authority jurisdictions and (e) an evaluation of our security infrastructure using AutoCore.**

I. Motivation & Introduction

Eyewitness accounts of an automotive accident are often inaccurate or conflicting. To mitigate this problem, we propose *AutoCore*, a VANET (Vehicular Ad hoc NETwork) application that automatically records video and telemetry data in a crash for use during an investigation. If investigators were provided with such data and could be assured of its authenticity, there would be less need to depend on eyewitness accounts. AutoCore could also help determine liability in hit-and-run incidents.

An automated collision reporting application presents an interesting set of research challenges, including maintaining vehicle location and identity privacy, providing conditional anonymity for vehicles reporting collisions, protecting the system against various attacks and ensuring the authenticity of the reported data.

Previous work that addresses VANET security and privacy has focused on identifying threats [15], [18], trusted inter-vehicle communication [9] and on the design of a security framework for VANETs [7], [11], [17]. Several challenges, however, remain open: First, the task of distributing the cryptographic credentials used by vehicles to sign and authenticate outgoing messages has been largely ignored. Second, proposed key management schemes require a centralized database for the conditional anonymity of vehicles [17], introducing a single point of failure. Third, a concrete example of how these proposed techniques could be applied to protect a particular VANET application has been missing.

Contributions – Our contributions are (a) the AutoCore crash reporting application and a detailed analysis of the threats against it, (b) a security infrastructure to protect AutoCore that uses a decentralized scheme to provide conditionally anonymous inter-vehicle communication, (c) the concept of cryptographically-verifiable Road-worthiness Certificates issued to vehicles by authorized vehicle service centers, (d) the use of these certificates in a practical scheme for the distribution of cryptographic vehicle credentials issued by regional authorities via roadside access points, (e) efficient support for the roaming of vehicles between different regions and (f) a detailed security and cost analysis of our infrastructure.

Because our focus is on security and privacy, a detailed description of AutoCore is beyond the scope of this paper. Instead, we provide a brief overview of the system in section III, after reviewing the state of the art in section II. Section IV analyzes our threat model, section V introduces our security infrastructure and section VI describes how we use this security infrastructure to protect AutoCore. Section VII analyzes our security infrastructure and section VIII discusses implementation issues. We review related work in section IX and briefly discuss some directions for future work when concluding in section X.

II. State of the Art

Automakers are working on pushing the safety envelope with proactive safety systems. These next-generation safety systems require vehicles to form cooperative groups, allowing them to exchange information and build awareness of their environments [8], [4], [23]. VANETs are thus a logical foundation for such safety systems.

To date, most industrial and academic research efforts in vehicular safety communications over VANETs have focused on the design of suitable MAC protocols, with the most prominent initiative being WAVE (Wireless Access for the Vehicular Environment, often also referred to as DSRC or Dedicated Short Range Communications). Designed as a short-to-medium range wireless protocol, WAVE provides data rates of up to 27Mbps over a range of 1000m and has been standardized as 802.11p [12]. In the US, the FCC has allocated radio spectrum for WAVE at the 5.9GHz band, with regulatory bodies in the EU and Japan pursuing similar initiatives. Examples

of vehicular safety applications studied so far include collision avoidance, cooperative driving and traffic optimization [4].

Vehicular safety communications primarily consist of independent *geocast* [1] messages produced by vehicles and roadside infrastructure. These messages typically fall into one of two groups [13]: *Routine Safety Messages* sent by vehicles and infrastructure on a regular basis, usually two or three times a second, and *Event Safety Messages* triggered by changes in vehicle behaviour, such as sudden braking, or infrastructure status, such as a vehicle running a traffic light. Messages generated by AutoCore fall into the latter category.

With VANETs largely being an emerging research field, little work has been done so far to address the security and privacy issues that arise from vehicles constantly sharing information about their movements and whereabouts with other vehicles and roadside infrastructure. One key challenge is the conditional anonymity of vehicles: a VANET security scheme should make it impossible for a global observer (e.g., law enforcement authorities or insurance companies) to track vehicles through the messages they transmit, while simultaneously allowing a vehicle to be reliably identified through these same messages when liability needs to be determined in the event of a crash and the ensuing investigation.

III. AutoCore

In this section, we present our automated collision reporting application. We begin by listing the concerned entities, then provide a brief overview of the system and end by walking through a typical usage scenario.

A. Concerned Entities

In designing AutoCore (and our supporting security infrastructure), we consider the following five groups:

1) Drivers: Drivers would allow the deployment of an application such as AutoCore in their vehicles only if they were given the right incentives (e.g., lower insurance premiums), if the application were completely automated and if it did not compromise their location or identity privacy in any way. (See Zimmer [24] for details about privacy issues in VANETs.)

2) Governmental Transportation Authorities: We assume that Governmental Transportation Authorities (GTAs) would be able to expand on their traditional roles as vehicle licensing authorities and operators of roadside infrastructure (traffic lights, stop signs, etc.) by issuing "smart" roadside infrastructure elements and vehicles with the cryptographic credentials required to secure inter-vehicle and vehicle-to-infrastructure communication.

3) Courts of Law: In the interest of privacy, drivers would want a trusted legal entity, ideally a Court of Law, to have control over the release of their identities as recorded in any crash evidence collected by AutoCore. Such a legal setup would be similar to that already in place for the tapping of phone lines and viewing of bank records (i.e., subpoenas).

[1] *Broadcast messages that contain information relevant only to recipients in a limited geographic region*

4) Law Enforcement Authorities: Law enforcement authorities would want easy access to authenticated evidence (video and vehicle telemetry) produced by AutoCore. We assume the presence of legal barriers, as described earlier, to prevent the abuse of this information by law enforcement authorities.

5) Roadside Access Point Operators: Operators of Roadside Access Points (RAPs) could be GTAs, law enforcement authorities or commercial service providers. We assume the deployment of RAPs at locations accessible by vehicles, such as fuel stations and public parking lots, with each RAP having Internet connectivity and potentially serving more than one purpose (e.g., the delivery of in-car entertainment content and electronic vehicle registration services).

B. System Overview

AutoCore consists of control software, secure storage and a software interface to on-board positioning, imaging and telemetry sensors. To support the system, we assume the presence of a Tamper-Proof Device (TPD), a positioning system such as Differential GPS, cameras such as those carried by luxury vehicles like the Lexus LS460 and the Mercedes-Benz S-Class for automated parking and enhanced driver night vision, respectively, and a WAVE-like communication interface for inter-vehicle and vehicle-to-infrastructure communication.

Vehicles with AutoCore continuously record video, temporarily storing this data with the corresponding vehicle positioning and telemetry data in their TPDs. This temporary storage takes the form of a ring buffer, where old data is overwritten by more recent data. In the event of a crash, AutoCore copies the last 60 seconds of data to persistent TPD storage as part of a collision report. The size of the persistent storage is determined by regional accident statistics and the average size of a collision report.

In this paper, we term vehicles that are directly involved in a crash *Primaries* and those in the vicinity (i.e., within camera range) *Witnesses*. Primaries and Witnesses equipped with AutoCore are capable of broadcasting two types of messages over their WAVE-like interfaces:

- *Collision Beacons*: to notify nearby vehicles that a collision event has occurred. Each beacon includes the current time and the source vehicle's GPS coordinates.
- *Witness Beacons*: to notify nearby vehicles that the source vehicle is a witness to a collision event. Each beacon includes the current time and the source vehicle's GPS coordinates.

In the event of a crash, Primaries equipped with AutoCore broadcast Collision Beacons, triggering the generation of collision reports by other Primaries and Witnesses (also equipped with AutoCore). Because vehicles continuously record video and positioning and telemetry data, AutoCore has access to data from before it finishes processing a Collision Beacon, ensuring processing and message propagation delays do not reduce the amount of useful data included in collision reports.

Law enforcement authorities may obtain collision reports from Primaries either by physically removing their TPDs to access the data they contain or by authenticating with AutoCore via handheld devices to obtain the data wirelessly through

the vehicles' WAVE-like interface. Similarly, reports produced by Witnesses may be obtained either via handheld devices or through delivery to RAPs, which forward these reports to law enforcement authorities over the Internet. We term RAPs providing report delivery services and handhelds issued to law enforcement authorities *Collision Report Collectors* (CRCs). Both types of CRCs are certified by the local GTA and hold similar cryptographic information.

For auditing purposes, when collision reports are delivered to a CRC or to law enforcement authorities, recipients would issue senders (vehicles and CRCs, respectively) with cryptographically verifiable receipts, as described in section VI-C.

Figure 1 shows the format of a typical collision report. These reports include all AutoCore messages pertaining to the incident and are signed to ensure their authenticity, with the public key certificate required for validation of this signature (see section V-F) also attached. The items shaded in dark gray are encrypted before delivery, as detailed in section VI-C.

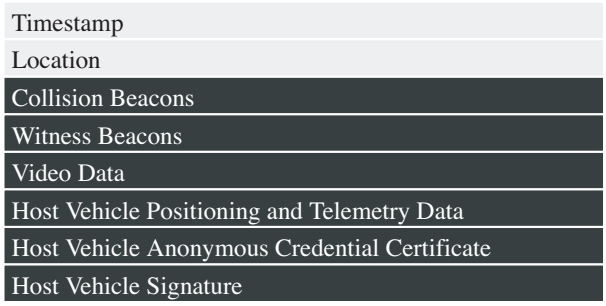

Fig. 1. AutoCore collision report.

When investigating a crash, law enforcement authorities can collate collision reports by the time and location of the incident in question, using the provided video, positioning and telemetry evidence to aid in reconstructing the crash. We describe how reports are decrypted and verified in section VI-C.

C. Usage Scenario

To clarify AutoCore's operation, we now walk through a usage scenario involving several vehicles fitted with the system travelling in opposite directions along a highway.

Vehicle A speeds up and attempts a late lane change, colliding with vehicle B. Both vehicles, equipped with AutoCore, are now termed Primaries. At the moment of impact, onboard sensors, such as those used to trigger airbags, inform AutoCore of a collision and both vehicles broadcast Collision Beacons. Vehicles C and D travelling behind the two Primaries as well as vehicle E traveling on the opposite side of the highway, but ahead of A and B, are equipped with AutoCore. These vehicles hear the Collision Beacons sent by one or both of the two Primaries. Vehicles C, D and E are now termed Witnesses. These Witnesses respond by broadcasting Witness Beacons, informing all involved vehicles (Primaries and Witnesses) of their presence. The vehicles record these Witness Beacons as well as the original Collision Beacons in their collision reports.

After the collision, vehicle A is severely damaged. Law enforcement officials obtain its collision report by removing its TPD. The reports produced by vehicles B and C (the latter stops after the collision) are obtained by law enforcement officials using handheld CRCs. Vehicles D and E automatically deliver their reports to roadside CRCs, which then forward the reports to law enforcement authorities over the Internet.

IV. Threat Model

In this section, we present a threat model for AutoCore that extends known threats against VANETs identified in related work [15], [18]. We begin by describing the potential capabilities of attackers and then move on to present general classes of threats against AutoCore.

In line with the general VANET attacker model presented by Raya and Hubaux [18], potential attackers considered in our threat model may be defined along four dimensions:

1) *Insider vs. Outsider* – insiders would possess valid credentials and be capable of abusing AutoCore protocols, whereas outsiders would be intruders without valid credentials and thus more limited capabilities.
2) *Malicious vs. Rational* – a malicious attacker or prankster would seek no personal benefit from attacking AutoCore, aiming only to harm users of the system, whereas a rational attacker might seek to subvert or disrupt the operation of AutoCore for personal benefit (e.g., to avoid liability in an accident).
3) *Active vs. Passive* – active attackers would be capable of eavesdropping on and generating AutoCore messages and reports, whereas passive attackers would be limited to eavesdropping on AutoCore communications.
4) *Local vs. Extended* – local attackers would be limited in scope to several nearby vehicles and/or RAPs, whereas extended attackers might have access to AutoCore entities distributed across a large geographical region.

Because it is impossible to identify every possible threat against AutoCore, we instead consider the following three general classes of threats:

- *False Information* – an attacker may broadcast invalid AutoCore messages to disrupt its operation or attempt to produce false collision reports. Similarly, in the event of malfunctioning or compromised equipment/software, AutoCore may produce authenticated messages and collision reports that contain false information.
- *Masquerading* – an attacker may masquerade as a legitimate AutoCore entity, generating data tagged with the legitimate entity's credentials to escape liability (in the case of a vehicle) or in order to intercept or modify valid collision reports (in the case of RAPs and CRCs).
- *Identity Disclosure and Tracking* – an attacker may use a combination of authenticated AutoCore entities (vehicles and RAPs) and base stations that eavesdrop on AutoCore communications to identify and potentially track the movements of vehicles equipped with the system.

A detailed discussion of how these classes of threats manifest themselves in AutoCore appears in the full-length version of this paper [21]. Other general classes of VANET threats, such as denial of service attacks and cheating with sensor

data, are described alongside potential solutions in related work [18].

V. Security Infrastructure

In this section, we begin by discussing the use of a TPD to store secret data and protect the integrity of AutoCore software in vehicles. We then describe the certificate authorities required to support our security infrastructure, introduce four types of cryptographic elements used by our infrastructure (Vehicle Identifiers, Road-worthiness Certificates, Electronic License Plates and Anonymous Credentials) and explain how each of them is generated and distributed. We end by describing an efficient scheme for the roaming of vehicles outside their home GTA's jurisdiction. A discussion of the efficiency of our proposed security infrastructure appears in section VIII-A.

A. Tamper-Proof Device

The protection of sensitive data stored in vehicles, such as collision reports produced by AutoCore and the cryptographic keys described in the following subsections, mandates a Tamper-Proof Device (TPD). We assume that a TPD is similar to a TPM (Trusted Platform Module), as defined by the Trusted Computing Group [20]. Namely, the device can generate key pairs and perform signing operations. Private keys never leave the device (or only in encrypted form). TPDs contain sensors that detect tampering and erase all the sensitive information protected by the device. With the help of the private keys embedded in the TPD, software using the TPD can authenticate a vehicle to roadside access points or to handhelds issued to law enforcement officials and prove that the TPD has not been tampered with, as described in sections V-E and V-F.

TPMs protect only against software-based attacks. Since VANETs are used for safety applications, we require TPDs to also resist hardware-based attacks. Furthermore, for software that uses a TPM, the TPM provides mechanisms to authenticate the state of this software at the software's load time. For TPDs, we assume that these mechanisms have been extended to ensure that the software is in a predefined state throughout its runtime and that sensitive information protected by the TPM become inaccessible as soon as this software is being tampered with. For example, secure co-processors provide this functionality by running software within a tamperproof box. However, secure co-processors tend to be expensive and slower than current desktop computers. The exact design of a TPD is therefore the topic of future research.

The TPD uses secure storage for storing collision reports and signing keys. This storage is either embedded in the device or external. In the latter case, the stored data must be encrypted, its integrity ensured and the TPD must defend against replay attacks.

In the rest of this paper, we have the term "TPD" cover both the actual TPD and any software that makes usage of the TPD and that is protected by the TPD, as explained above.

B. Certificate Authorities

We envision the presence of several certificate authorities to support our security infrastructure:

1) Vehicle Manufacturers: Currently, vehicle manufacturers issue a unique Vehicle Identifier Number (VIN) to all vehicles that they produce. These numbers are stamped onto the frame of a vehicle, effectively binding VINs to vehicles for their operational lifetime. Similarly, manufacturers will issue vehicles with *Vehicle Identifiers* (see section V-C) that can be cryptographically verified and are bound to the vehicle for its operational lifetime.

2) Governmental Transportation Authorities (GTAs): Just as these authorities register vehicles and issue physical license plates, GTAs will issue *Electronic Licence Plates* (see section V-E) to vehicles registered in their region of jurisdiction. In addition, GTAs will issue vehicles that operate in their jurisdiction and that hold valid Electronic Licence Plates (not necessarily issued by the same GTA) with *Anonymous Credentials* (see section V-F), which allow the vehicles to communicate with other vehicles in the region. We assume that GTAs have certificates for all vehicle manufacturers registered in their jurisdiction.

C. Vehicle Identifiers

Vehicle Identifiers are used to uniquely identify vehicles. A vehicle identifier consists of a signing key pair (VID_{Pu}, VID_{Pr}) and a corresponding certificate VID_{Cert}, which contains information uniquely identifying the vehicle (e.g., its VIN number) and that binds this information to the vehicle's public key VID_{Pu}. Such a public key is equivalent to today's VIN numbers. The pair (VID_{Pu}, VID_{Pr}) is created by a vehicle's TPD and VID_{Cert} is issued and installed in a vehicle's TPD by its manufacturer during production. Vehicle Identifiers are valid for the lifetime of a vehicle.

D. Road-worthiness Certificates

A Road-worthiness Certificate $RoadWorthy_{Cert}$ is issued to a vehicle by its manufacturer or by authorized inspection authorities. Such a certificate proves that the vehicle has been inspected and approved for road-worthiness (safety checks, emissions, etc.). The certificate lists the vehicle's VID_{Pu} and is valid for the period of time for which the vehicle has been deemed to be road-worthy. We assume that the vehicle's home GTA holds certificates for each inspection authority, which allows the GTA to validate Road-worthiness Certificates. These certificates are used by a vehicle to renew its Electronic License Plate.

E. Electronic License Plates

Electronic License Plates (ELPs) serve the same purpose as physical license plates. ELPs consist of a signing key pair (ELP_{Pu}, ELP_{Pr}) and a corresponding certificate ELP_{Cert}, which binds the vehicle's VID_{Pu}, contained in VID_{Cert}, to its public key ELP_{Pu} under a digital signature produced by the vehicle's home GTA. The certificate is valid for the duration of the vehicle's registration (about a year).

Vehicles initially acquire or renew their ELPs through roadside access points (RAPs) using the Road-worthiness Certificates described in section V-D. Figure 2 shows the renewal protocol. The protocol works as follows:

When a vehicle reaches a RAP that advertises ELP renewal services for its home GTA, the vehicle's TPD authenticates

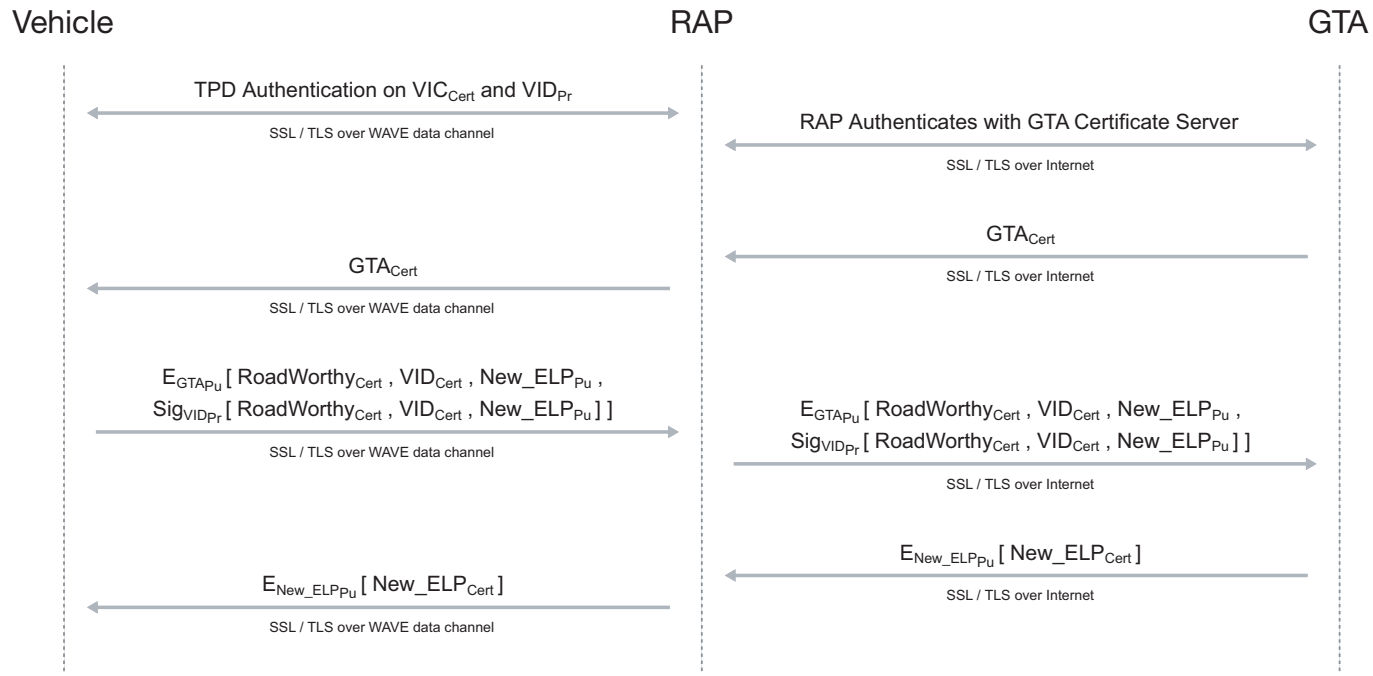

Fig. 2. Protocol for the renewal of ELPs.

with the RAP using its VID_{Cert} and VID_{Pr}. For example, the authentication can use SSL/TLS with client authentication. The purpose of client authentication is to demonstrate to the RAP that the vehicle has not been tampered with; any tampering with the TPD would have erased VID_{Pr}, preventing the vehicle from authenticating. After successful authentication, the RAP will be willing to act as a relay between the vehicle and the local GTA.

To guard against compromised RAPs and attackers masquerading as legitimate RAPs, each RAP has a signing key pair (INF_{Pu}, INF_{Pr}) and a corresponding certificate INF_{Cert}, issued by the local GTA. The TPD ensures the validity of this certificate when authenticating with the RAP.

The vehicle first receives the public key of the GTA, signed by a publicly known CA, such as VeriSign. We assume that the TPD has some CA certificates embedded in it, similar to a Web browser, and information that allows it to identify certificates belonging to GTAs.

The TPD then generates a signing key pair (New_ELP_{Pu}, New_ELP_{Pr}) and sends New_ELP_{Pu}, along with $RoadWorthy_{Cert}$ and VID_{Cert}, to the RAP for forwarding to the GTA. The message is signed with VID_{Pr}. For privacy reasons, the message is encrypted with the public key of the GTA. Once the GTA has decrypted the ciphertext and verified the certificates, the vehicle's signature and potentially other, external conditions, such as payment of fees or traffic tickets, it issues New_ELP_{Cert} covering New_ELP_{Pu} to the vehicle via the RAP. Similar to the vehicle-RAP connection, the RAP-GTA connection is also secured with SSL/TLS.

F. Anonymous Credentials

Anonymous Credentials consist of a signing key pair ($AnonCred_{Pu}$, $AnonCred_{Pr}$) and a certificate $AnonCred_{Cert}$ covering $AnonCred_{Pu}$. The certificate is issued by a GTA (not necessarily a vehicle's home GTA) and contains no public information that could be used by an unauthorized observer to identify the vehicle. Vehicles will possess a set of Anonymous Credentials and use the signing key $AnonCred_{Pr}$ of a credential to sign outgoing AutoCore messages. The corresponding certificate $AnonCred_{Cert}$ accompanies such a message. To avoid tracking of a vehicle based on $AnonCred_{Cert}$, the vehicle changes credentials often using a variable-frequency key changing algorithm [17].

A certificate $AnonCred_{Cert}$ consists of

$AnonCred_{Pu}$ | $InvisibleIdentity$ | GTA_GUID | $Sig_{GTA_{Pr}}$[$AnonCred_{Pu}$ | $InvisibleIdentity$ | GTA_GUID].

The $InvisibleIdentity$ field in the certificate consists of

$E_{Court_{Pu}}$ [$E_{GTA_{Pu}}$ [VID]].

GTA_GUID is a globally unique identifier assigned to the issuing GTA and VID is a unique identifier (per GTA) assigned to each ELP (similar to a license plate number). The two identifiers are included in an ELP. $AnonCred_{Cert}$ includes GTA_GUID in plaintext so that the identity of the GTA (and the court) that can decrypt the $InvisibleIdentity$ field can be determined.

An $InvisibleIdentity$ is invisible because it is first encrypted using the issuing GTA's public key and then encrypted again using the public key of a trusted legal entity (in our case, a local Court of Law). This double encryption ensures that a vehicle's identity is hidden and can be revealed only when both the local court of law and the GTA co-operate. Note that we need a probabilistic encryption scheme for producing the $InvisibleIdentity$ field. This way, a vehicle's $InvisibleIdentity$ field will be different in each of its Anony-

mous Credentials, preventing tracking of the vehicle based on this field.

To ensure the conditional anonymity of vehicles (as described in section II), we use a blind signature scheme [6] for the certification of Anonymous Credentials by GTAs. Our scheme has the advantage that a GTA cannot learn a vehicle's $AnonCred_{Pu}$'s while being ensured that the vehicle's identity can be recovered from $InvisibleIdentity$ (if approved by a court). We present the protocol in figure 3. We now briefly outline the details of this protocol.

When a vehicle encounters a RAP that advertises Anonymous Credentials refresh services, it authenticates with the RAP using the same process described earlier for the renewal of ELPs. In the next step, the vehicle gets the certificates for the local GTA, the local court and the local law enforcement authorities. The certificates are all signed by a publicly known CA. The vehicle also gets certificates for neighbouring GTAs of the local GTA. These certificates will be used when the vehicle travels between GTA jurisdictions (see section V-G). The vehicle then executes the following protocol:

1. First, the vehicle computes the number of Anonymous Credentials that it will require given the amount of credentials that it already holds and its distance-to-empty (remaining driving range given current fuel). This number, N, can be computed using the variable-frequency key-changing algorithm mentioned earlier.

2. The vehicle's TPD then generates N key pairs ($AnonCred_{Pu}$, $AnonCred_{Pr}$) and produces certificate templates for each of these ($TmAnonCred_{Cert}$). The templates contain all the information contained in a certificate, except the GTA's signature. The templates are then blinded and sent to the GTA, along with the vehicle's electronic license plate ELP_{Cert}. This information is encrypted with GTA_{Pu} and signed with ELP_{Pr}.

3. The GTA decrypts the ciphertext, validates the signature and ensures that ELP_{Cert} has not expired. If successful, the GTA signs each of the blinded certificate templates and returns them to the vehicle.

4. The vehicle then unblinds the signatures and combines them with the original templates.

In this protocol, the GTA signs blinded certificate templates, preventing it from verifying whether these templates have the required structure, as shown above. Furthermore, it cannot check whether the value encrypted in the $InvisibleIdentity$ field is correct. Because certificate templates are generated by the TPD and the vehicle's request message is signed with its ELP_{Pr}, the GTA can assume that the TPD has not been tampered with and that the template and the embedded ciphertext are correct. To deal with malfunctioning TPDs that, for example, include a wrong VID in the $InvisibleIdentify$ field, we could use a cut-and-choose protocol, where the vehicle sends $M > N$ blinded templates to the GTA, which in turn asks the vehicle to unblind $M - N$ randomly chosen templates before signing the remaining ones. The main disadvantage of using such a cut-and-choose protocol is the increased load on a GTA, particularly when $M \gg N$.

G. Travel between GTA Jurisdictions

Vehicles communicating under our security scheme, as described so far, are only capable of authenticating (and thus reacting to) messages generated by vehicles and infrastructure from their home GTAs. Therefore, we need to extend our scheme to allow vehicles to communicate with vehicles and infrastructure certified by other GTAs, enabling communication while traveling outside the home GTA's jurisdiction.

Before proposing our solution, we introduce some terminology. Vehicles operating outside their home GTA's jurisdiction are termed *Visitors*, while vehicles/infrastructure operating/deployed within their home GTA's jurisdiction are termed *Locals*. A *Foreign GTA* is the GTA responsible for the region in which a *Visitor* is operating, while a *Home GTA* is the GTA responsible for the region the Visitor is registered in.

In our solution, we allow Visitors to acquire Anonymous Credentials from a Foreign GTA. This approach requires that the Foreign GTA maintains a list of trusted GTAs, one of which being the Visitor's Home GTA. This way, the Foreign GTA can verify ELPs certified by the Home GTA. Assuming a Visitor will ultimately encounter a RAP that provides Anonymous Credentials refresh services, this approach guarantees that the Visitor will ultimately be able to communicate with other vehicles and infrastructure in the Foreign GTA's region.

In the interval before a Visitor encounters this RAP in the Foreign GTA's region, it will not be able to communicate with other vehicles or infrastructure. We propose the following solution for this problem: as described in section V-F, vehicles receive a set of $NeighbourGTA_{Cert}$ certificates when they obtain Anonymous Credentials. This set includes GTA certificates for each neighbouring Foreign GTA, $ForeignGTA_{Cert}$'s, signed by the Home GTA, and a set of certificates for the Home GTA, $HomeGTA_{Cert}$'s, each signed by a different neighbouring Foreign GTA. Similarly, each infrastructure element, like a RAP or a handheld CRC, is given the same set of certificates when it is certified by its Home GTA. We assume cooperation between neighbouring GTAs to achieve this.

With this technique, Visitors are able to authenticate messages produced by Locals by verifying the certificate attached to such a message against the set of $ForeignGTA_{Cert}$'s that the Visitors hold.

To allow Locals to authenticate Visitors' messages, Visitors include one of their $HomeGTA_{Cert}$'s in their messages, in particular, the certificate issued by the Foreign GTA. Locals validate $HomeGTA_{Cert}$ and add the public key of the Visitor's Home GTA to their list of trusted GTAs.

Including $HomeGTA_{Cert}$ in messages increases overall message size (see section VIII-A for actual message sizes). To reduce communication channel congestion, a Visitor could add $HomeGTA_{Cert}$ to a subset of its messages. We leave a more thorough exploration of this technique to future work. Note that, as soon as a Visitor obtains Anonymous Credentials from the Foreign GTA, it no longer has to transmit $HomeGTA_{Cert}$.

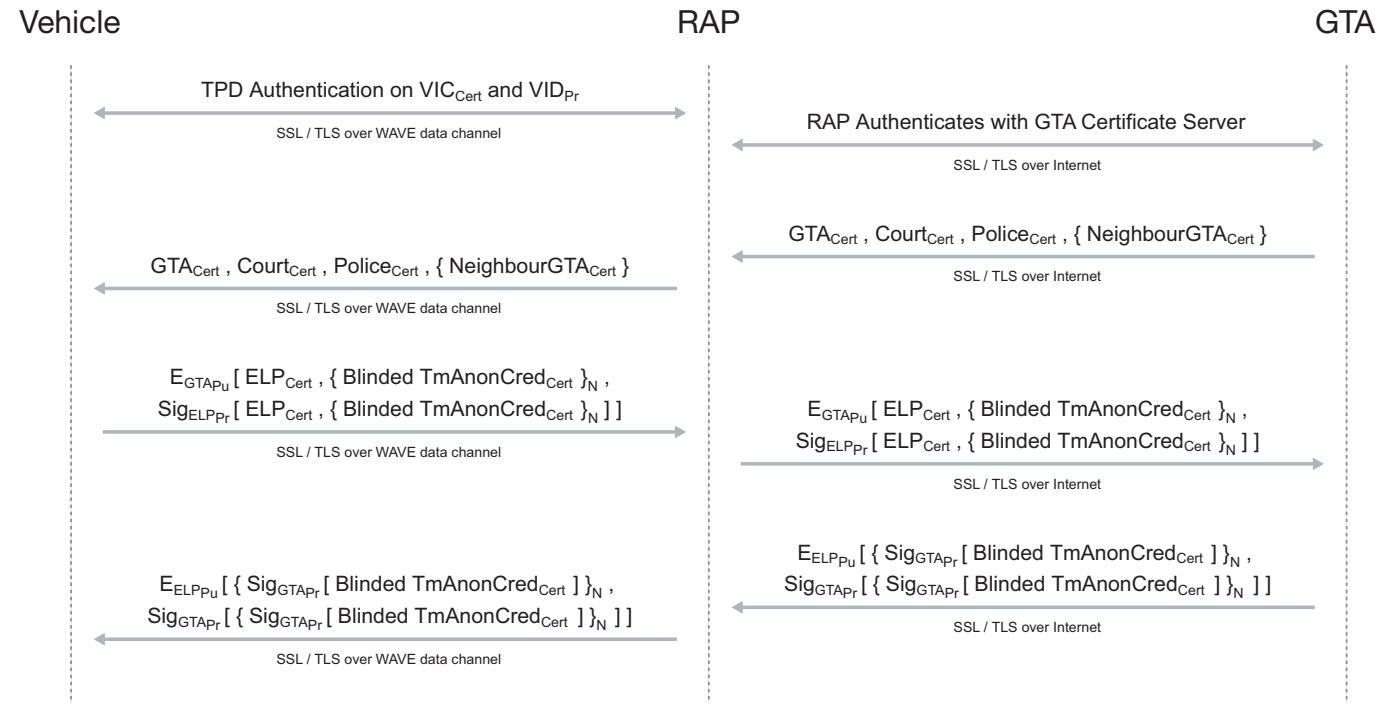

Fig. 3. Protocol for the renewal of Anonymous Credentials.

VI. Securing AutoCore

In this section, we discuss how the security infrastructure introduced earlier is used to protect AutoCore communications and collision reports, as described in section III-B.

A. Securing Inter-Vehicle Communications

In securing inter-vehicle communications, we are primarily concerned with ensuring the authenticity of AutoCore messages and guaranteeing non-repudiation of these messages while protecting vehicle location and identity privacy. We achieve these goals through the use of Anonymous Credentials.

Collision and Witness Beacons produced by AutoCore follow the format shown below:

$$M, T, Sig_{AnonCred_{Pr}}[M|T], AnonCred_{Cert}.$$

M is the message, T is a timestamp included to ensure message freshness and $AnonCred_{Cert}$ is the Anonymous Credential certificate signed by the local GTA that corresponds to $AnonCred_{Pr}$ used to sign the message. The signature guarantees message authenticity and non-repudiation. If necessary, a court and a GTA can jointly determine the identity of the sending vehicle by decrypting the $InvisibleIdentity$ field in $AnonCred_{Cert}$. We discuss this process in section VII-C.

B. Securing Vehicle to Infrastructure Communications

Communication between vehicles and infrastructure is required for the delivery of data. Examples of such data include the cryptographic credentials held by vehicles, as described in section V, and collision reports produced by AutoCore. Vehicle to infrastructure communication is secured using standard mutual authentication and secure data transfer protocols, such as SSL/TLS with client authentication.

As mentioned in section V-E, we assume that each RAP has a signing key pair (INF_{Pu}, INF_{Pr}) and a corresponding certificate INF_{Cert}. We make the same assumption for other kinds of roadside infrastructure, such as handhelds issued to law enforcement officials. This way, vehicles are able to detect fake infrastructure.

C. Securing AutoCore Collision Reports

As mentioned in section IV, it is necessary to cryptographically protect collision reports in order to guarantee their integrity and prevent the abuse of information they contain.

To describe how reports are secured, we refer back to the AutoCore collision report format shown in figure 1. When vehicles have finished recording video evidence and telemetry data, they place this data in a report with all Collision and Witness Beacons received for the corresponding collision event. Each report's header contains the event's timestamp and location, as recorded by the vehicle generating the report. The report is then signed with the reporting vehicle's current private key $AnonCred_{Pr}$, with the corresponding certificate $AnonCred_{Cert}$ included for verification purposes. These items combine to produce the complete report shown in figure 1. The items shaded in dark gray are encrypted using the public key $Police_{Pu}$ of the local law enforcement authority when a vehicle encounters a CRC to deliver its report to. Vehicles obtain this public key in $Police_{Cert}$ when obtaining Anonymous Credentials. Reports that cannot be retrieved from a vehicle through a CRC, perhaps because the vehicle is badly damaged, are obtained directly from the vehicle in cleartext form, as described in section III-B.

When law enforcement authorities obtain these reports, they decrypt them using $Police_{Pr}$, if necessary, and verify the

submitting vehicle's signature to ensure that the report is authentic. Next, they can revoke the conditional anonymity of the Beacon messages based on the process discussed in section VII-C. Note that the police should decrypt a collision report only if certain standards have been met (e.g., the accident or the hit-and-run incident have been reported). In particular, for privacy reasons, the police should not proactively decrypt collision reports.

During each delivery step, the CRC or law enforcement authority receiving a collision report returns a signature, created with INF_{Pr} or $Police_{Pr}$ and covering the received report, to the delivering vehicle or RAP. This process establishes a cryptographically-verifiable audit trail.

VII. Security Analysis

In the following subsections, we describe how the use of our security infrastructure guards against the three general classes of threats outlined earlier in section IV.

A. False Information

All messages produced by legitimate vehicles are signed with their current $AnonCred_{Pr}$ and authenticated using the corresponding $AnonCred_{Cert}$. While this only guarantees that a message comes from a vehicle that was trustworthy when it was issued the $AnonCred_{Cert}$, it does prevent outsiders (as defined in section IV) from sending authenticated messages. When a vehicle's TPD is tampered with, it automatically erases all its $AnonCred_{Pr}$'s and other private keys, preventing the vehicle from generating authenticated messages or further renewing its credentials.

In the event of an AutoCore equipment malfunction or successful tampering with the system, we depend on verifying data consistency. In AutoCore, data consistency is provided by having AutoCore correlate received messages against similar messages produced by other vehicles in close time and space. Messages could also be correlated with host vehicle telemetry data. For example, sudden braking might indicate that the driver is reacting to an emergency. In addition, if more than one vehicle equipped with AutoCore is involved in a collision, it is possible to correlate messages sent by these vehicles by checking the timestamps and locations included in the messages. Golle et al. [10] discuss techniques for data correlation in VANETs.

B. Masquerading

We prevent masquerading and provide non-repudiation in vehicles as follows:

- A vehicle cannot claim to be a different vehicle, because it signs messages with its own private keys. Furthermore, ELPs are unique and only one vehicle holds the corresponding ELP_{Pr} in its TPD.
- A vehicle cannot deny having sent messages because a message is signed with $AnonCred_{Pr}$, which belongs to the vehicle and was generated by the vehicle in the first place. Timestamps included in each message guard against message replay attacks.

Similarly, legitimate RAPs and CRCs use their unique INF_{Pr} keys to authenticate with vehicles and law enforcement authorities when forwarding AutoCore data, issuing receipts and processing credentials refresh requests.

C. Identity Disclosure and Tracking

To guard against the disclosure of vehicles' identities and the tracking of their movements from the messages they broadcast, these messages are made anonymous. This anonymity is conditional in that in the event of an accident, authorized entities can revoke this anonymity and identify vehicles from the messages they sent. We now address these two issues in more detail.

1) Conditional Anonymity: Privacy against vehicle identification attacks is guaranteed through the absence of any public information about a vehicle in $AnonCred_{Cert}$'s sent out by the vehicle. Furthermore, a GTA that issues an $AnonCred_{Cert}$ does not see the contents of these certificates and will thus not be able to re-identify vehicles by colluding with roadside infrastructure.

Privacy against vehicle tracking attacks is guaranteed through frequently changing the Anonymous Credential used by a vehicle. Since each Anonymous Credential looks different, these credentials cannot be used for tracking a vehicle. Furthermore, a vehicle must also frequently change its MAC and IP addresses.

It should be noted, however, that as with other proposed VANET security schemes [18], our scheme is effective only at preventing the tracking of vehicles when two or more vehicles equipped with the system are operating nearby. This is similar to the concept of k-anonymity [19], in that the larger the number of vehicles using the system in an area, the higher the privacy guarantees for these vehicles.

2) Revocation of Conditional Anonymity: In the event of a collision, law enforcement authorities may want to learn the identities of vehicles that sent messages included in collision reports. To reveal these identities, the authorities will take a report to a Court of Law. For each beacon message in the report, the court will remove the first layer of encryption from the $InvisibleIdentity$ field in $AnonCred_{Cert}$ attached to the message. The second layer of encryption will be removed by the GTA upon receipt of a valid court order, which will reveal a vehicle's identity.

VIII. Implementation Issues

In this section, we examine two implementation issues, namely, what kind of cryptosystem to use in our security infrastructure and how to implement imaging in AutoCore.

A. Choice of Cryptosystems

Every safety message sent out by a vehicle, such as a Collision Beacon, contains, in addition to its payload, a digital signature and a certificate for the corresponding public key. To reduce overhead, we need cryptosystems with short signature and key sizes. We choose ECDSA for most of the signing key pairs. The only exception is the signing key pair used by GTAs for certifying Anonymous Credentials. As mentioned in section V-F, these signatures are issued in a blind way.

However, no blind signature scheme based on ECDSA is currently known. Instead, we use a blind signature scheme based on the BLS short signature scheme [1], referred to as BBLS ("Blind BLS") in the rest of this section. We describe this scheme in the full-length version of this paper [21].

For ECDSA, we choose a public key size of 20 bytes, which results in signatures of 40 bytes. This setup is secure in the short and medium term, which is sufficient for our purposes since there is no long-term storage of safety messages. For BBLS, we choose a public key size of 75 bytes, which results in signatures of 25 bytes. Note that BBLS public keys are not exchanged in safety messages.

Finally, we choose the Elliptic Curve Integrated Encryption Scheme (ECIES) [5] for encrypting a car's identity in the $InvisibleIdentity$ field. ECIES also generates a MAC, which we drop because the integrity of the $InvisibleIdentity$ field is assured by the certificate in which the field is embedded. Using 20 byte public keys, the size of the $InvisibleIdentity$ field will be 44 bytes, that is, 4 bytes for the actual ciphertext (we assume that VID has a size of 4 bytes) and 20 bytes for each of the two random elliptic curve points (i.e., their x-coordinates) output in the two encryption steps.

Given this setup, the security overhead of a safety message is 133 bytes. In particular, the ECDSA signature of the message is 40 bytes long. The accompanying certificate has a size of 93 bytes: 20 bytes for $AnonCred_{Pu}$, which is an ECDSA public key, 25 bytes for the BBLS signature, 4 bytes for the GTA_GUID and 44 bytes for the $InvisibleIdentity$ field. For comparison, Xu et al. [22] estimate the typical payload size of a safety message to be between 100 and 400 bytes. The overhead in Raya and Hubaux's scheme [18], which relies on 28 bytes ECDSA keys, is 140 bytes.

As mentioned in section V-G, when traveling between GTA jurisdictions, a vehicle might temporarily include a certificate in which the Foreign GTA certifies the vehicle's Home GTA in some of its messages. The size of such a certificate is 60 bytes; the public key of the Home GTA needs 20 bytes and the ECDSA signature requires 40 bytes.

A vehicle generating a safety message needs to create an ECDSA signature. Similarly, a vehicle receiving a safety message needs to verify this signature and the BBLS signature of the certificate accompanying the signature. On a Pentium IV 3 GHz, it takes about 0.68 ms to generate an ECDSA signature and 1.3 ms to verify the signature. It takes about 49.7 ms to verify a BBLS signature (and 2.8 ms to create it). Raya and Hubaux [17] estimate that a vehicle has only about 2.5 ms for processing a message, which is much less than the time it takes to verify a BBLS signature. However, this signature needs to be validated only once for an Anonymous Credential used by a vehicle. Any additional messages using the same credential no longer require this overhead. For example, assume the same scenario as introduced by Raya and Hubaux [17], a highway with six lanes (three in each direction) and an inter-vehicle distance of 30 m. Vehicles transmit safety messages every 300 ms over a 300 m communication range. Here, within a second, a vehicle driving at 120 km/hour is going to see about nine new Anonymous Credentials from vehicles traveling in the opposite direction and six from vehicles traveling in the same direction, which leaves sufficient time to check the signature of each new Anonymous Credential.

B. AutoCore Imaging

For imaging data, precise requirements will likely vary depending on the jurisdiction a vehicle is operating in (i.e., the evidence requirements of local law enforcement authorities), but a very minimal setup consists of an omni-directional video camera mounted on the roof or boot of a vehicle. One such system developed by Peri and Nayar [16] is capable of taking images from an omni-directional camera and generating pure perspective images.

IX. Related Work

Gerlach [9] highlights key VANET security concepts and proposes a model for trusted inter-vehicle communication. Parno and Perrig [15] further examine VANET security issues, identify potential attacks and introduce a categorization scheme for adversaries. We build on these threats in section IV.

Hubaux et al. [11] focus their efforts on vehicle identity and location privacy, introducing Electronic License Plates (ELPs) that serve the same purpose as physical license plates. In Hubaux et al.'s scheme, an ELP is simply an identifier. Instead, the ELPs proposed in our work consist of signing key pairs and certificates. This way, it becomes possible to use ELPs to refresh Anonymous Credentials via RAPs. Hubaux et al. also introduce the concept of Electronic Chassis Numbers (ECNs) that can be used to uniquely identify vehicles. Our Vehicle Identifiers serve a similar purpose, but instead of simply being identifiers, they also consist of signing key pairs and certificates and can be used to bootstrap the renewal of ELPs via RAPs. Note that although Hubaux et al. mention that ELPs can be renewed upon vehicle registration, they do not present an actual renewal scheme.

Raya and Hubaux [17] build on this earlier work and introduce a security framework for VANETs, proposing the use of *Anonymous Keys* to sign messages sent by vehicles. To provide conditional anonymity, Anonymous Keys mandate the creation of a centralized database that maps these keys to a vehicle's identity. This approach has the drawback that the database becomes a single point of failure. To avoid abuse of this database, Raya and Hubaux suggest encrypting this database with shared secrets split between authorities. Raya and Hubaux do not elaborate on the certification and distribution process of Anonymous Keys. In particular, there is the danger that, while these keys are being certified, a malicious certificate issuer, like a GTA (or an intruder), can store mappings between public keys and identities in a second database, which is not protected with a secret-sharing system. In our approach, as explained in section V-F, the certifier never sees the public key in the Anonymous Credential that is being certified. Furthermore, our approach includes the information necessary for revoking anonymity directly in Anonymous Credentials, thereby eliminating the need for a

centralized database. Finally, we describe an actual scheme for the generation and distribution of Anonymous Credentials.

To reduce the amount of time a malfunctioning or rogue vehicle can communicate with other vehicles, Jungels et al. [14] introduce the RTPD and RCCRL protocols. Another protocol proposed by the same authors, DRP, can be used in pure ad hoc mode, with vehicles accumulating accusations against misbehaving vehicles and reporting these to the GTA via a RAP. Key problems with these CRL schemes are their pervasive infrastructure and vehicle tracking requirements.

In order to allow vehicles to communicate with infrastructure and other vehicles in regions governed by foreign GTAs, Raya and Hubaux [18] suggest the use of base stations deployed at borders to verify and re-certify a vehicle's Anonymous Keys. The disadvantage of this approach is that it requires that a (working) base station be deployed at every single border crossing or even just a crossing between two provinces/states, in case the provinces/states are governed by different GTAs (as in the case of Canada or the US). This assumption seems unrealistic. Moreover, when vehicles do not stop when crossing an inter-GTA border, the amount of time required to verify and to re-certify a vehicle's Anonymous Keys may be too large, with moving vehicles going out of range before the process is completed. (WAVE is currently limited to a range of 1 km [12].) Our solution avoids these problems through the use of $NeighbourGTA_{Cert}$ certificates, eliminating the need for base stations at borders.

Choi et al. [7] present a VANET security architecture in which only access points, but not cars, authenticate messages sent by cars. Authentication exploits MACs. This architecture is not suitable for AutoCore. Namely, AutoCore messages (and other Event Safety Messages) are sent only in safety-critical situations, which makes message authentication by a receiving car mandatory. Furthermore, collision reports can be used for determining liability in an accident, which requires non-repudiation, a feature not offered by MACs.

Instead of Anonymous Credentials, we could also use an anonymous credential system, such as Idemix [3] or Brands credentials [2]. However, these systems tend to be more expensive. For example, issuing a Brands credential takes three steps, whereas our protocol requires only two.

X. Conclusion

We have introduced AutoCore, an automated crash reporting application that provides cryptographically-verifiable evidence of an automobile crash in the form of digital video and telemetry data recorded by vehicles either involved in or at the scene of the crash. To secure AutoCore, we have presented a security infrastructure that extends the state of the art in VANET security. We have analyzed this security infrastructure to demonstrate its robustness and efficiency.

Directions for future work include the design of an efficient CRL scheme for Anonymous Credentials and an evaluation of our security infrastructure to assess its suitability for other VANET safety applications.

Acknowledgments

We are grateful to S. Keshav for suggesting the crash reporting application and for precious comments on a preliminary version of this paper. We also thank Ian Goldberg for suggesting and extending the BLS short signature scheme [1] and for suggesting the ECIES encryption scheme [5].

References

[1] D. Boneh, H. Shacham, and B. Lynn. Short Signatures from the Weil Pairing. *Journal of Cryptology*, 17(4):297–319, 2004.

[2] S. Brands. A Technical Overview of Digital Credentials. Technical report, Credentica, February 2002.

[3] J. Camenisch and E. Van Herreweghen. The Design and Implementation of the Idemix Anonymous Credential System. In *Proceedings of 9th ACM Conference on Computer and Communications Security (CCS 2002)*, pages 21–30, November 2002.

[4] Crash Avoidance Metrics Partnership (CAMP). *Vehicle Safety Communication Final Report*. 2006.

[5] Certicom Research. Standards for Efficient Cryptography, SEC 1: Elliptic Curve Cryptography, Version 1, September 2000.

[6] D. Chaum. Blind Signatures for Untraceable Payments. In *Proceedings of CRYPTO '82*, pages 199–203, 1982.

[7] J.Y. Choi, M. Jakobsson, and S. Wetzel. Balancing Auditability and Privacy in Vehicular Networks. In *Proceedings of ACM Q2SWinet '05*, October 2005.

[8] W. Enkelmann. FleetNet - Applications for Inter-Vehicle Communication. In *Proceedings of the IEEE Intelligent Vehicles Symposium '03*, pages 162–167, June 2003.

[9] M. Gerlach. VaneSe - An Approach to VANET Security. In *Proceedings of V2VCOM 2005*, July 2005.

[10] P. Golle, D. Greene, and J. Staddon. Detecting and Correcting Malicious Data in VANETs. In *Proceedings of 1st Workshop on Vehicular Ad Hoc Networks*, October 2004.

[11] J. P. Hubaux, S. Capkun, and J. Luo. The Security and Privacy of Smart Vehicles. *IEEE Security and Privacy Magazine*, 2(3):49–55, 2004.

[12] *IEEE P1609.2 Version 1 – Standard for Wireless Access in Vehicular Environments: Security Services for Applications and Management Messages*. 2006.

[13] D. Jiang, V. Taliwal, A. Meier, W. Holfelder, and R. Herrtwich. Design of 5.9 Ghz DSRC-based Vehicular Safety Communication. *IEEE Wireless Communications Magazine*, 13(5):36–43, 2006.

[14] D. Jungels, M. Raya, I. Aad, and J. P. Hubaux. *Certificate Revocation in Vehicular Ad Hoc Networks*. Technical Report LCA-REPORT-2006-006. EPFL, 2006.

[15] B. Parno and A. Perrig. Challenges in Securing Vehicular Networks. In *Proceedings of 4th Workshop on Hot Topics in Networks (HotNets-IV)*, November 2005.

[16] V. N. Peri and S. K. Nayar. Generation of Perspective and Panoramic Video from Omnidirectional Video. In *DARPA Image Understanding Workshop (IUW)*, pages 243–246, May 1997.

[17] M. Raya and J. P. Hubaux. The Security of Vehicular Ad Hoc Networks. In *Proceedings of 3rd ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN 2005)*, pages 11–21, November 2005.

[18] M. Raya and J. P. Hubaux. Securing Vehicular Ad Hoc Networks. *Journal of Computer Security*, 15(1):39–68, February 2007.

[19] L. Sweeney. k-Anonymity: a model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 10(5):557–570, 2002.

[20] Trusted Computing Group. https://www.trustedcomputinggroup.org. Accessed February 2007.

[21] S. Ur Rahman and U. Hengartner. *Secure Crash Reporting in Vehicular Ad hoc Networks*. Technical Report CACR 2007-11. Centre for Applied Cryptographic Research, University of Waterloo, 2007.

[22] Q. Xu, T. Mak, and R. Sengupta. Vehicle-to-Vehicle Safety Messaging in DSRC. In *Proceedings of 1st International Workshop on Vehicular Ad Hoc Networks*, pages 19–28, October 2004.

[23] X. Yang, J. Liu, F. Zhao, and N. Vaidya. A Vehicle-to-Vehicle Communication Protocol for Cooperative Collision Warning. In *Proceedings of MobiQuitous '04*, 2004.

[24] M. Zimmer. Personal Information and the Design of Vehicle Safety Communication Technologies. In *AAAS Science & Technology in Society Graduate Conference*, April 2005.