

Security & Privacy for DSRC-based Automotive Collision Reporting

Sumair Ur Rahman, Hossein Falaki
DRC School of Computer Science, University of Waterloo, Ontario, Canada
{surrahman, mhfalaki}@cs.uwaterloo.ca

ABSTRACT

In this paper, we present the security and privacy features of a DSRC-based application we designed to provide authenticated eyewitness information about an automobile crash in the form of digital photographs and telemetry captured during an emergency event by vehicles involved in or in the vicinity of a collision.

1. MOTIVATION & INTRODUCTION

For those of us who commute by car on a regular basis, seeing the aftermath of a crash is unfortunately an accepted reality. Although emergency services usually respond to these incidents quickly, removing injured persons within minutes of a crash, wreckage often remains on the scene for several hours whilst accident investigators (usually the Police) take photographs and measurements in an effort to determine liability. For other commuters travelling on the same route, this often means lengthy delays; as traffic has to be diverted around the site and people slow down to take a look.

In order to help solve this problem, we set about designing a DSRC-based application that automatically recorded photographic and telemetry data during an emergency event for use during a crash investigation. If investigators were provided with such data and could be assured of its authenticity, wreckage at crash sites could be removed quickly, reducing traffic delays. In addition to serving this purpose, such a system could also help determine liability in hit-and-run incidents. From a commercial perspective, such a system would likely also prove attractive to insurance companies. If it were easier and less costly to determine liability in the event of an accident, they would have incentive to

lower insurance premiums, which of course is good news for drivers.

From a research perspective, an automated collision reporting application presents an interesting set of 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 reported data.

Since our focus is on security and privacy, a detailed description of such a system is beyond the scope of this paper. Instead, we provide a brief overview of how such a system might work in section 3. Section 4 then analyses our threat model, bringing us to section 5 where we introduce the security model we developed to support this and other DSRC-based safety applications. Section 6 provides a brief feasibility study, particularly in terms of message sizes and data storage costs. Before concluding, we mention some potential future work in section 7.

2. STATE OF THE ART

Vehicular Communications (VC) is an emerging field with only a few pioneering papers in the areas of security and privacy. Gerlach introduces security concepts for vehicular networks in [10]. In [7], Hubaux et al. focus on identity and location privacy, introducing the concept of Electronic License Plates (ELPs) that serve as unique identifiers for vehicles. Parno and Perrig classify adversary types and discuss potential attacks on vehicular networks in [4]. In [3], Raya and Hubaux describe a security framework for VANETs, proposing the use of multiple anonymous keys to sign messages sent by vehicles in order to ensure location and identity privacy. In [1], Raya et al. compare the security mechanisms required by VC with those required by more conventional networks, demonstrating the need for, and opportunity to use, novel security solutions.

3. SYSTEM MODEL

In this section we provide an overview of DSRC, briefly describe typical safety messages and then introduce AutoCore, our hypothetical automotive collision reporting system.

3.1. DSRC Overview

Dedicated Short-Range Communications (DSRC) consists of a short- to medium-range wireless protocol specifically designed for vehicle-to-vehicle and vehicle-to-infrastructure communication. Based on the 802.11a standard, DSRC is in the process of being standardized as 802.11p by the IEEE 1609 working group.

In the US, DSRC has been allocated a 75MHz wide spectrum (7 x 10Mhz licensed channels) around 5.9GHz by the FCC, with Europe and Japan operating a similar sized spectrum at 5.8GHz. Designed for both low overhead and latency, DSRC typically delivers data at speeds between 6 and 27 Mbps over a range of 1000m (given line of sight).

Whilst current-generation (915MHz) DSRC applications primarily consist of electronic toll collection, international programs such as Intelligent Transportation Systems (ITS) are bringing together governments, industry and academia with the goal of utilizing technologies such as next-generation 5.9GHz DSRC to reduce traffic congestion and improve safety through VC applications such as crash prevention, cooperative driving and traffic flow optimization. Several prototype vehicles and applications were recently demonstrated at the 2006 ITS World Congress in London, UK [12].

3.2. Safety Messaging Fundamentals

As described in [2], safety messages exchanged between vehicles and roadside infrastructure fall into one of the two categories described below.

Routine Safety Messages – these are status messages sent by vehicles on a regular basis, usually two or three times a second. Such messages only remain useful for a short period of time, typically a few seconds and are mainly intended to allow other vehicles on the same road (as well as roadside infrastructure such as traffic lights) to predict the vehicle's movements.

Event Safety Messages – these are either triggered by changes in vehicle behaviour (such as sudden braking) or infrastructure status (such as a

vehicle running a traffic light) that break the continuity implied by routine safety messages. All messages generated by our collision reporting system fall into this category.

As part of a simulation and analysis of their proposed framework in [3], Raya and Hubaux estimate that a vehicle travelling in dense highway traffic could expect to receive up to 400 such safety messages per second. In terms of messages sizes, the CAMP project determined (as part of its study of potential safety applications in [11]) that most safety messages would be smaller than 100 bytes, with only a few being larger (up to 430 bytes for the left turn assistance application).

3.3. AutoCore

In this subsection, we describe AutoCore, our hypothetical collision reporting application. We begin by listing the concerned entities, then provide a typical usage scenario and finally describe the messages produced by our system. Security constructs specified in these messages will be discussed in detail in subsequent sections.

3.3.1. Concerned Entities

In the design of AutoCore, we considered the needs and roles of the following five parties.

Drivers – despite operating vehicles fitted with the system, drivers would likely not want to have to actively submit collision reports (the system has to be able to do this automatically, without any driver involvement), or have the compromise their location or identity privacy (see [13] for a detailed discussion of privacy issues in VSC). It goes without saying that the AutoCore system should not be a distraction on the road.

Vehicle Manufacturers – since most safety applications and the associated hardware will likely be OEM, manufacturers shipping these systems with new models would primarily be concerned with minimizing costs. As such, the system would have to utilize as many hardware components already available in cars or used by other safety applications as possible. We assume the presence of a tamper-proof device as described in [3] and the availability of cameras (these are slowly making their way in luxury vehicles such as the Lexus LS460 and Mercedes-Benz S-Class to support automated parking and enhanced driver night-vision, respectively).

Transportation Authorities – primarily concerned with maintaining roadside infrastructure (traffic lights, stop signs, etc) and issuing license

plates, transportation authorities would want minimal involvement beyond their traditional role as a licensing authority in the collision reporting process.

Certificate Authorities – currently only concerned with issuing certificates for PKI, given the incentive of a new revenue stream, CAs could expand their service offerings to include issuing specialized asymmetric keys and certificates for vehicles to their manufacturers. With their expertise in securing storing data, we also assume that CAs would be willing to store these keys and certificates for investigative purposes (only to be released upon a court order, protecting the privacy of drivers).

Investigators – we include the Police and insurance companies in this category as entities who would want to investigate and determine liability in the event of an accident. Given only the time and location of an incident, these parties would want to have easy access to authenticated evidence (photographs and vehicle telemetry). We assume that the requirement of a court order to obtain such evidence is not unrealistic.

3.3.2. Usage Scenario

Before describing a typical use case, we introduce some terminology. Vehicles directly involved in a collision are termed *Primaries* whereas those in the vicinity and within camera range (usually about 200m) are termed *Witnesses*. Vehicles equipped with the AutoCore system are capable of sending four types of messages over DSRC:

- *EmergencyStartBeacon* – informs nearby vehicles of an imminent emergency event, giving this event a unique *EmergencyEventID* (composed of a random number, the current timestamp and GPS location).
- *CollisionBeacon* – informs nearby vehicles that a collision has occurred in the emergency event with the specified *EmergencyEventID*.
- *WitnessBeacon* – informs nearby vehicles that the source vehicle is a witness to the emergency event with the supplied *EmergencyEventID*.
- *EmergencyEndBeacon* – informs nearby vehicles that the emergency event with supplied *EmergencyEventID* has ended.

As described in [3], we assume that no two emergency events share the same *EmergencyEventID* and that AutoCore runs and stores its data in a tamper-proof device (we call it a *Collision Event Recorder*, or *CER*) alongside other safety applications. For the sake of simplicity, we assume that DSRC communication is reliable and that

Witnesses are within communication range of *Primaries* when an emergency event ends. It should be noted that neither of these assumptions necessarily hold in practice. For the delivery of Witness Reports (photos and telemetry recorded by Witnesses), we assume the presence of roadside access points operated by the local Transportation Authority (TA) called *Collision Evidence Collectors (CECs)* with secure connections to a server running at the TA.

We now walk through a usage scenario involving several vehicles fitted with the system travelling in both directions on a highway.

Vehicle A speeds up and attempts a late lane change, coming dangerously close to vehicle B. Sensing the danger, the driver of vehicle B reacts with sudden braking and change of direction, while onboard safety systems issue a collision warning. This warning is also delivered to AutoCore, which reacts by broadcasting an *EmergencyStartBeacon*, signalling a potential collision. Vehicle B begins recording data (photos and telemetry) as part of a *Collision Report*.

Nearby vehicles fitted with the system hear the *EmergencyStartBeacon* and respond by broadcasting *WitnessBeacon* messages to announce their presence as witnesses to the incident. *Primaries* and *Witnesses* hear these messages and include them as part of their reports, creating a list of all vehicles present at the scene. Like vehicle B, they too begin recording photos and telemetry as part of a *Witness Report* for submission to roadside access points in the event of a collision.

The drivers of both *Primaries* fail to avert a collision and onboard sensors detect impact. The signals produced by these sensors on vehicle B are heard by AutoCore, which now broadcasts a *CollisionBeacon*, confirming the occurrence of a collision. *Witnesses* hear this beacon and know that the emergency event has in fact resulted in a collision; their *Witness Reports* should be delivered to a roadside access point.

When *Primaries* come to rest, sensors alert AutoCore that stops capturing data and broadcasts the *EmergencyEndBeacon*. *Witnesses* stop recording data and photos when either of the following events occurs: an *EmergencyEndBeacon* message is received; the *Witness* has travelled a certain distance (if it was moving) or a certain time has elapsed (if it was stationary). If *Witnesses* do not receive a *CollisionBeacon* before this, all recorded data and photos are discarded (AutoCore determines that a potential collision was averted).

Collision and Witness Reports are authenticated and stored securely in the vehicles' *CERs*. Accident investigators obtain a *Collision Report* from vehicle B by pulling out its *CER* (the

tamper-proof device functions as a black box), whilst *Witness Reports* are automatically delivered to CECs installed at gas stations and major intersections as *Witnesses* pass by them.

To accommodate for the fact that an entire report may not be delivered while a *Witness* is in range of a *CEC*, we utilize a delay-tolerant connection with the CEC, such as that described by Seth et al. [14]. To secure such a connection, we would utilize techniques similar to those described by Seth and Keshav in [15]. It is also worth noting that CECs may not necessarily have always-on connectivity to the TA server, in fact, in the case of highway installations, there may not be any connectivity at all. In such situations, we would use mechanical backhaul via authorized vehicles operated by law enforcement and/or the TA to securely transport data from such CECs to the TA's server. This would be similar to the scheme described by Keshav et al. [17].

Reports received by the TA server are stored in a database keyed by *EmergencyEventID*. Investigators simply need to provide this unique identifier (or specify an incident's date, time and location to search the database) to obtain all reports submitted for an incident.

To track the progress of reports and provide Witnesses, CECs and vehicles providing mechanical backhaul between CECs and the TA server with proof-of-delivery (an audit trail), each is issued a receipt when it successfully forwards a report to the CEC, a vehicle providing mechanical backhaul or the TA server respectively. Since all *Witnesses* are identified in Collision and Witness reports submitted to the authorities, any *Witnesses* withholding evidence can be identified if necessary.

3.3.3. Application Messages

A detailed description of the messages broadcast by AutoCore described in the previous section appears in the appendix to this paper.

3.3.4. Hardware Support

To ensure the integrity of the AutoCore system, it would be required to run on a tamper-proof device as described in [3]. Instead of interfacing directly with on-board sensors to detect potential collisions (emergency events), AutoCore would likely rely on other on-board safety applications such as Electronic Stability Control (ESC) for this information.

For imaging data, precise requirements would likely vary depending on the jurisdiction the

vehicle is operating in, but a very minimal setup might consist of an omni-directional video camera mounted on the roof or boot of a car. One such system developed by Nayar and Peri described in [18] and [19] is capable of taking images from an omni-directional camera and generating pure perspective images. The adaptation of such computer vision systems for use in bad weather and low-light conditions is described in [20].

4. THREAT MODEL

In this section we present our threat model, identifying potential attacks against the AutoCore system and categorizing them as either threats to privacy or system security.

4.1. Privacy Threats

Our privacy goals for the AutoCore system are to provide conditional anonymity such that vehicles at the scene of an accident can't be tracked or identified through the application messages they broadcast, whilst providing a means to identify *Primaries* and *Witnesses*, should accident investigators require such information to enforce liability.

4.1.1. Vehicle Positioning

As illustrated in [4], an attacker may attempt to track the movements of a vehicle by listening to the messages it broadcasts in order to learn about the driver. Although AutoCore only broadcasts messages when it encounters an emergency event (not very often), it would likely share keys with other safety applications housed in the same TPD that broadcast routine safety messages (up to 3 times a second). As such, our security scheme has to take into account the privacy threats posed by such applications using the same keys. The most significant threat in this case is made possible through the pervasive deployment of DSRC-capable roadside infrastructure. So-called smart traffic lights, dividers and signs could report received messages back to a central server, allow the operator to track the movement of vehicles passing by these installations.

4.1.2. Vehicle Identification

Whilst drivers already give up a significant amount of privacy by driving cars with license plates on

them, what they do expect is to only be identified when they are within visual range (i.e. while their license plates can be seen). As such, the ability to determine the identity of a vehicle without seeing it is considered a violation of privacy. An example of this would be roadside infrastructure such as traffic lights being able to identify vehicles from the messages these vehicles broadcast.

4.1.3. Leaked Collision/Witness Report Data

If not protected, it is possible for Witness and/or Collisions reports to be viewed by unauthorized persons. An insurance company, for example, might look through these reports to single out drivers who have been at the scenes of multiple accidents and charge them higher premiums. As such, it would be considered a privacy violation if accident investigators (the Police or insurance companies), other drivers, the Transportation Authority, or the Certificate Authority could freely view these reports. Clearly, a mechanism is required by which allows legal barriers to prevent the abuse of this information.

4.2. Security Threats

In this section, we consider threats to the authenticity of the data generated by AutoCore, the ability of attackers to subvert the system, attacks that would affect the integrity of AutoCore software and DoS attacks that affect system availability.

4.2.1. Denial of Service

Several types of Denial of Service (DoS) attacks are possible on this system. The most basic type of attack mentioned in almost all VANET security literature is signal jamming. In such an attack, an adversary would simply jam the communication channel used by vehicles and/or access points, rendering the application useless and preventing critical information from reaching other vehicles and access points. This paper does not try to address this type of denial of service attack. Several solutions are proposed to this type of attack in [3], including frequency hopping and channel switching.

An adversary could overwhelm vehicles by flooding them with false application beacons, rendering the communication channel and dependent applications useless. A similar type of DoS attack might target roadside access points used to deliver collision reports generated by witness vehicles.

A third type of attack might attempt to overwhelm the transportation authority by patching into the link between roadside access points and the transportation authority, flooding the channel with garbage data. This threat would likely have a similar solution as that to the previous attack.

4.2.2. Message Suppression

In this type of attack a driver either physically disables his inter-vehicle communication system or modifies the application to prevent it from either sending or responding to application beacons. The goal of such an attacker would be to prevent registration and insurance authorities from learning about collisions involving his vehicle and/or to avoid delivering collision reports to roadside access points.

4.2.3. Message Fabrication/Alteration

A prankster might fabricate or replay altered messages to force on-scene vehicles into recording collision data. These fictitious emergency/collision events would result in bogus data being collected by vehicles and making its way up to transportation authority, wasting valuable communication, processing and storage resources. An attacker might also use this technique to mask the occurrence of a collision by diverting limited application resources to the fabricated collision.

4.2.4. Key/Certificate Replication

In this type of attack, an adversary would seek to undermine the system by replicating a single vehicle's identity across several vehicles. The goal of such an attacker would be to confuse the authorities and possibly prevent identification (if the attack were carried out on a large enough scale) of vehicles in hit-and-run incidents.

5. SECURITY MODEL

Our security goals are as follows. We intend to guarantee the authenticity and privacy of all collision/witness reports, ensure the integrity and availability of the AutoCore system and provide conditional anonymity for all vehicles using the system. Further, there should be no way for the concerned authorities (investigators, transportation and certificate authorities) to abuse this conditional anonymity. To achieve the goals, we propose the use

of two classes of asymmetric keys. We now briefly describe the basic elements of our security model, the roles played by the concerned entities introduced earlier and the security protocols used by our system.

5.1. Security Elements

In this subsection we describe the two classes of asymmetric keys used by our security model, the entities that issue them, the entities that hold them and how they are installed. We also briefly describe our use of Anonymous ELPs (AELPs), a variation on the ELP concept introduced by Hubaux et al. in [7].

5.1.1. ECN

An Electronic Chassis Number (ECN) is another concept introduced by Hubaux et al. in [7]. Installed by a vehicle manufacturer in each vehicle's TPD before it rolls out of the factory, these signed certificates uniquely identify vehicles by their physical chassis numbers and are valid as long as the TPD has not been tampered with (tampering with it should erase/corrupt the ECN). A vehicle never discloses its ECN to another entity.

5.1.2. VID Keychains

All vehicles in our model hold a *VID Keychain*, each of which consists of one *VIDpub Key*, one *VIDsigning Key-pair* and N *VID Anonymous Certificates*. These "anonymous certificates" are different from standard certificates in that they do not contain a public key (technically, we probably shouldn't even be calling them certificates); instead they contain a random identifier called a *VID CertificateID* that only the issuing CA can use to look up the corresponding private key. The use of "anonymous certificates" ensures only the CA is capable of decrypting data encrypted with *VIDpub* and that entities observing messages containing these "anonymous certificates" cannot learn the sender's *VIDpub* key. We leave the calculation of N (how many such "anonymous certificates" each vehicle needs) to future work, although one should be able to expand on a similar computation that appears in [3].

VID Keychains are valid for a year and are refreshed each time the vehicle goes in for its annual check-up. CAs issue *VID Keychains* to vehicle manufacturers who then use an Anonymous Credentials scheme similar to the IDEMix system described in [9] on the vehicle's ECN to install the keychain. The use of such a scheme to install *VID*

Keychains ensures the manufacturer cannot track which vehicle has been issued which keychain.

5.1.3. Anonymous ELP

We build on the Electronic Licence Plate (ELP) idea introduced in [7] to create Anonymous ELPs. Unlike ELPs, which can be used directly identify a vehicle; an AELP includes a random identifier that maps to the vehicle's real ELP (with this mapping known only to the transportation authority). The other information contained by an AELP is that same (vehicle registration region, dates of validity, etc) as that included in an ELP. Currently, we envisage each vehicle holding one AELP for each day its registration is valid for (so usually 365).

5.1.4. COMM Keychains

The idea of using multiple keys to sign messages generated by vehicles appears in several papers. Using each key in only one stretch of road (see [3] for an example of how the length of this stretch of road can be estimated) ensures that the vehicle remains anonymous. We call the combination of such a communication key and certificate a *COMM Keychain*. M such keychains are issued to each vehicle by the Transportation Authority (TA) responsible for the jurisdiction it is registered in. The issuing TA signs certificates in each *COMM Keychain*. For an example of how M can be determined, please refer to [3].

COMM Keychains are installed by a vehicle's local TA using the same sort of Anonymous Credentials scheme used by vehicle manufacturers. In this case however, the credential used by the vehicle to authenticate itself with the TA is its current AELP. As such, when a vehicle comes in for its annual registration, it is first issued its set of AELPs and then its *COMM Keychains*.

CECs and other roadside infrastructure are also issued *COMM Keychains* so that they may authenticate with passing vehicles to collect collision/witness reports and exchange data. To guard against key compromise, the TA would have to refresh these keys on a regular basis.

5.2. Concerned Entities

In this subsection we view our security model from the perspective of the entities in our system model to further clarify the use of the security elements introduced in the previous subsection.

5.2.1. Vehicles

Each vehicle holds all of the security elements introduced earlier. We summarize these for the sake of completeness:

- 1 x *ECN*: issued by the vehicle's manufacturer and never revealed by the vehicle to another entity. Valid for the lifetime of the vehicle and used by the vehicle to bootstrap *VID Keychains*.
- 1 x *VID Keychain*: used by the vehicle to hide its identity in the messages it broadcasts and reports it submits to roadside access points. The use of N "anonymous certificates" ensures the vehicle can't be tracked through its use of only one certificate for its *VID Keychain*.
- 365 x *AELPs*: used to identify the vehicle as part of our conditional anonymity scheme. One *AELP* is valid for each day of the year.
- *M* x *COMM Keychains*: used to sign messages broadcast by the car and authenticate with roadside access points (or other infrastructure).

5.2.2. Roadside Access Points

Each roadside access point is issued a single *COMM Keychain* by the TA operating it. These keychains are periodically refreshed to guard against key compromise.

5.2.3. Transportation Authorities

Expanding on their traditional roles as vehicle registration authorities, Transportation Authorities (TAs) are responsible for issuing and installing *AELPs* and *COMM Keychains* in vehicles. The use of an Anonymous Credentials system to deliver the *COMM Keychains* (bootstrapped on the vehicle's current *AELP*) ensures TAs are not able to keep track of which vehicle has been issued which keychain. TAs maintain a database mapping *AELPs* to *ELPs*. As mentioned earlier, TAs are also responsible for operating roadside APs and issuing the same with *COMM Keychains*.

5.2.4. Certificate Authorities

Certificate Authorities (CAs) expand on their traditional roles by issuing *VID Keychains* to vehicle manufacturers for installation in vehicles. Since CAs

do not install the keychains themselves, they cannot keep track of which vehicle is using which keychain.

As mentioned in section 5.1.2, CAs maintain a database of all issued keychains, mapping each keychain to its corresponding *VID CertificateIDs*. When an entity investigating an incident (the Police or an insurance company) wishes to decrypt data secured by a *VIDpub Key* (only the CA knows the corresponding *VIDpri Key*), it must present the CA with a court order requiring it to re-encrypt the data for the investigator. In this case, a proxy re-encryption scheme (similar to that described in [8]) is used to re-encrypt the data for the investigator, ensuring the *VIDpri Key* remains secret (any other data encrypted with it is not revealed to the investigator).

5.3. Security Protocol

In this subsection the security protocol used for vehicle-to-vehicle communication and vehicle to access point communication.

5.3.1. Access Point Communication

Communication between vehicles and APs, between APs and mechanical backhaul vehicles, and between APs or the mechanical backhaul vehicles and the TA server are secured by the STS protocol [21]. Entities use their *COMM Keychains* to engage in a handshaking protocol that both mutual authenticates both parties and allows them to establish a shared secret to encrypt and sign transferred data.

5.3.2. Inter-Vehicle Communication

To guard against message replay/fabrication and DoS attacks, every valid message needs to be authenticated. Two levels of authentication are utilized for AutoCore messages. The structure of a typical message can be seen in figure 2. Noteworthy features include:

- Every message contains the *AELP* of the vehicle and the hash of the public portion of the *COMM Key* (that was used to sign the whole message). This information is encrypted and signed with the vehicle's *VID Keychain*. The current *VID Anonymous Certificate* is also attached.
- Each message is signed with the vehicle's current *COMM Key*, with the corresponding *COMM Key Certificate* attached so receiving vehicles can verify the signature.

5.3.3. Securing Collision/Witness Reports

Data contained in Collision/Witness reports generated by colliding and witness vehicles is encrypted and signed using a random symmetric key. This symmetric key in turn, is encrypted along with other meta information on a *VIDpub Key*. The vehicle's current *VID Anonymous Certificate* is attached to the message before it is signed using the vehicle's current COMM Keychain. The incident's timestamp, location and *EmergencyEventID* appear in plaintext at the top of the report. As such, the contents of reports are kept secret from the drivers of vehicles, roadside access points, the transportation authority and even the investigators.

To find all reports for specific event, investigators need to supply the TA (which stores all the reports) with a timestamp and location or *EmergencyEventID* (known if the investigators pulled a CER from a damaged vehicle).

To view the contents of a report, investigators must check the *VID Anonymous Certificate* to determine which CA issued the certificate. With a court order in hand to investigate the event, the investigators then approach the CA to proxy re-encrypt the report so they can read it. At this point, the CA issues them a special key to decrypt the report and the re-encrypted report, keeping the *VIDpri Key* used to secure other messages/reports sent by the vehicle secret.

5.3.4. Conditional Anonymity

The presence of an *AELP* in the message enables conditional anonymity. To determine the identity of a vehicle that sent a message, investigators follow a similar process to that described in the previous section when a court issues them permission to do so. Since the *AELP* is also anonymous, investigators must contact the transportation authority indicated in the *AELP* to determine the true identity of the vehicle and its owner. This last check ensures the CA is unable to abuse its position of holding all the *VIDpri Keys* for the *VID Keychains* it has issued.

6. FEASIBILITY

In this section we will analyze the feasibility of deploying the proposed system, particularly in terms of message size and disk storage capacities.

One of the most important limitations of the DSRC event safety based applications is the size of messages. The very high frequency of message

exchange events (400 times per second) makes message size an important concern. In addition, sufficient integrity checking information is required in every message to guard against the security threats described in section 4.

Providing conditional anonymity requires encrypting data with public keys and attaching the corresponding certificate to the messages. All this extra information increases the size of the messages. With a typical DSRC safety message being less than 100 bytes in length and certificates taking an additional 350 bytes each, we expect the size of a typical message generated by our system to be around 800 bytes. This is 300 bytes over the total message sizes described in [3]. Clearly, our choice of cryptosystem depends on both encryption over-head (how much encrypted data is larger than the corresponding plaintext), the corresponding certificate's size and the time required to perform signature verification (identified as the bottleneck in VSC security mechanisms by Hubaux et al. in [3]).

The issue of storage space vs message size came up several times during the design of our security model. Put simply, disk storage is cheap. According to [22], in 2004 we could purchase 8.7MB of storage for 1 cent. CAs don't need to store entire *VID Anonymous Certificates*, only the corresponding *CertificateIDs*, *VIDpri Keys* and *VIDsigning_pub Keys*. So, if each car has 50'000 such anonymous certificates (see [3]) and each CertificateID is 16 bytes long (and that's really big), each vehicle requires the CA to allocate 0.76MB of storage. That means we can store all the VID CertificateIDs for 11 vehicles for less than 1 cent.

7. FUTURE WORK

Potential future work would include a simulation of our security model to determine its feasibility and the best cryptosystem to use (Hubaux et al. identify ECC as an ideal candidate in [3]). In addition, it would also be interesting to see how big our messages actually get (by this we mean total size; the application data in our messages rarely exceeds 100 bytes) once we have selected a cryptosystem.

Another potential expansion of this work might be determining how many *VID Anonymous Certificates* a vehicle needs to carry and investigating whether VID Keychains can be shared amongst several cars, further increasing anonymity.

8. CONCLUSIONS

In this paper we have presented a security model for a DSRC-based collision reporting application that requires guaranteed authenticity of reports, a verifiable audit trail and conditional anonymity of all participating vehicles. In addition, our security model also guards against known VANET attacks such as message fabrication/replay to ensure system availability and integrity.

Since reports submitted by vehicles contain data that can be abused by the authorities handling them, our security model was also designed to incorporate legal checks and balances such as court orders to protect the location and identity privacy of participating vehicles and their owners. Our conditional privacy scheme ensures that vehicles can be accurately identified for liability purposes if the required legal obligations have been met, without compromising the privacy of other reports and messages generated by the same vehicle.

Our use of an anonymous credentials scheme prevents authorities and vehicle manufacturers from tracking which keys are issued to which vehicles.

9. REFERENCES

- [1] M. Raya, P. Papadimitratos and J.P. Hubaux, "Securing Vehicular Communications," *IEEE Wireless Communications Magazine*, vol. 13, no. 5, October-November 2006, pp. 8-15.
- [2] 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*, vol. 13, no. 5, October-November 2006, pp. 36-43.
- [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] B. Parno and A. Perrig, "Challenges in Securing Vehicular Networks," *Wksp. Hot Topics in Networks (HotNets-IV)*, 2005.
- [5] P. Golle, D. Greene and J. Staddon, "Detecting and Correcting Malicious Data in VANETs," *Wksp. Vehicular Ad Hoc Networks (VANET)*, 2004.
- [6] "IEEE P1609.2 Version 1 – Standard for Wireless Access in Vehicular Environments: Security Services for Applications and Management Messages," in development, 2006.
- [7] J.P. Hubaux, S. Capkun and J. Luo, "The Security and Privacy of Smart Vehicles," *IEEE Security and Privacy Magazine*, vol. 2, no. 3, May-June 2004, pp. 49-55.
- [8] G. Ateniese, K. Fu, M. Green and S. Hohenberger, "Improved Proxy Re-Encryption Schemes with Applications to Secure Distributed Storage", in proceedings of *NDSS '05*.
- [9] J. Camenisch and A. Lysyanskaya, "Design and Implementation of the IDEMix Anonymous Credential System", in proceedings of *ACM CCS '02*.
- [10] M. Gerlach, "VaneSe: An Approach to VANET Security," *V2VCOM*, 2005.
- [11] Crash Avoidance Metric Partnership, "Vehicle Safety Communication Final Report," available through the U.S. Department of Transportation.
- [12] 2006 Intelligent Transportation Systems World Congress, www.itsworldcongress.com
- [13] M. Zimmer, "Personal Information and the Design of Vehicle Safety Communication Technologies: An Application of Privacy as Contextual Integrity," in the *Science & Technology Society Conference*, April 2005.
- [14] A. Seth, S. Bhattacharyya, and S. Keshav, "Application Support for Opportunistic Communication on Multiple Wireless Networks," *Manuscript*, November 2005.
- [15] A. Seth and S. Keshav, "Practical Security for Disconnected Nodes," in proceedings of *First Workshop on Secure Network Protocols (NPSEC)*, November 2005.
- [16] A. Seth, P. Darragh, S. Liang, Y. Lin, and S. Keshav, "An Architecture for Tetherless Communication," *Manuscript*, July 2005.
- [17] A. Seth, D. Kroeker, M. Zaharia, S. Guo and S. Keshav, "Low-cost Communication for Rural Internet Kiosks Using Mechanical Backhaul," in proceedings of *MOBICOM 2006*, Sept. 2006.
- [18] S. K. Nayar, "Omnidirectional Video Camera," *DARPA Image Understanding Workshop (IUW)*, pp. 235-242, May 1997.
- [19] V. N. Peri and S. K. Nayar, "Generation of Perspective and Panoramic Video from Omnidirectional Video," *DARPA Image Understanding Workshop (IUW)*, pp. 243-246, May 1997.
- [20] S. K. Nayar and S. G. Narasimhan, "Vision in Bad Weather," in proceedings of *IEEE International Conference on Computer Vision (ICCV)*, vol. 2, pp. 820-827, 1999.
- [21] A.J. Menzes, A.J. Vanstone and P.C.V. Ooschot, "Handbook of Applied Cryptography," CRC Press, 1996.
- [22] Historical Notes about the cost of HD Storage Space, www.alts.net/ns1625/winchest.html