Misbehavior Detection in V2X Communications

Kevin Henry

ESCRYPT, Waterloo, Ontario, Canada

Abstract

Connected vehicle technology consists of Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communication, collectively referred to as V2X. This enables vehicles and infrastructure to exchange safety related information to enable smarter, safer roads. If driver alerts are raised or automated action is taken as a result of these messages, it is critical that messages are trustworthy and reliable. To this end, the Security Credential Management System (SCMS) has been proposed to provide authentication and authorization of V2X messages without compromising individual privacy. A critical aspect of this system is the ability to identify and remove misbehaving devices from the network. This paper provides an overview the SCMS, proposed approaches to misbehavior management (or lack thereof), and some of the difficulties the SCMS is likely to encounter as it is more widely deployed.

Keywords:

Security, V2X, Misbehavior, Security

1 Introduction

Secure Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communication, collectively referred to as V2X, allow vehicles and infrastructure to communicate with each other, thereby enabling a new generation of safety and intelligent mobility solutions. By sharing sensor data with others on the road, smarter decisions can be made by all, and unsafe situation and be predicted and avoided. The United States Department of Transportation (USDOT) has predicted that this technology can reduce or eliminate up to 80% of non-alcohol related collisions.

One of the primary V2X safety applications currently proposed is the broadcast of Basic Safety Messages (BSMs), which are beacons containing basic telemetry information about a vehicle. This includes data such as the speed, heading, acceleration, brake status, and optional information about the vehicles past movement and expected future movement. Other road users can use these messages to, for example, detect a vehicle on a collision course, detect a

vehicle approaching a red light without braking, or to assist with vehicles making a left turn when their view is obstructed by another vehicle.

Infrastructure, such as traffic lights, may also receive BSMs and update their behavior to optimize the flow of traffic based on received messages. Approaching vehicles can be detected at a distance, and standard messages can help guide vehicles through the intersection more efficiently. Signal Phase and Timing (SPaT) messages contain information about the current phase of the lights, and MAP messages provide information detailing intersection movements, such as dedicated turning lanes. Additional infrastructure, such as smart signs, can transmit various roadside alerts to help inform vehicles and their drivers of their surroundings.

While V2X provides connected vehicles with a plethora of information that can be used to make smarter decisions, whether automated or not, the system only has value if the information is accurate and trustworthy. Traditional mechanisms for providing this trust do so by binding information to a verifiable identity. Naively applying such an approach to V2X messaging would allow each V2X message to be tied to a specific user, creating an easy way to track and surveil participants in the network. To avoid this problem, a new, privacy-preserving public-key infrastructure has been proposed to secure V2X communications. The Security Credential Management System (SCMS) was purpose-built to solve the problem of authenticating V2X messages without sacrificing individual privacy.

The SCMS has seen a great deal of attention, is deployed in many pilot demonstrations, and is poised to support national-scale deployments in the future. Despite this, a critical portion of the system remains relatively undeveloped: Misbehavior detection. If a vehicle is observed transmitting incorrect data, whether intentionally or due to misconfiguration or malfunction, then it must be actively removed from the network until the cause of the misbehavior is identified and addressed. Misbehavior detection is further complicated by the private nature of the system. Devices must be revocable without compromising privacy.

This remainder of this paper presents a brief introduction to the technical aspects of the SCMS and provide an overview of the current state of misbehavior detection and weaknesses in the current proposals.

2 The Security Credential Management System

The SCMS is a large-scale Public Key Infrastructure (PKI) with several novel enhancements intended to specifically support connected vehicles and roadside equipment. The system issues *credentials* which allow communicating devices to trust only authorized sources. The security technology is built around ECC, a highly efficient public key algorithm that can operate at high

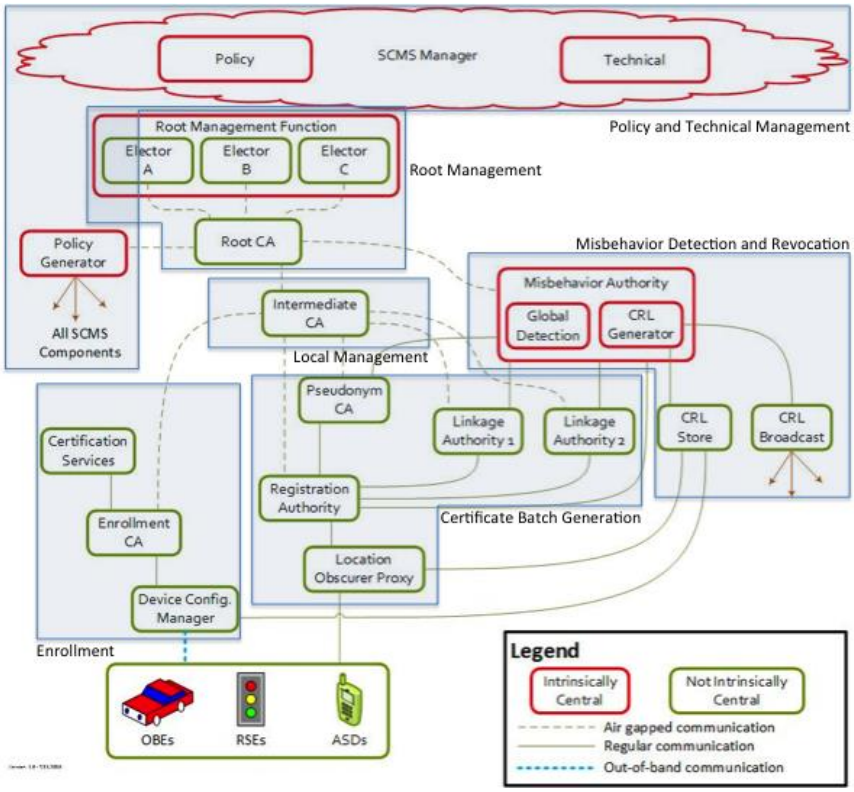

Figure 1. Annotated SCMS Architecture Diagram.

speed on embedded processors and enables small digital signatures. It takes advantage of implicit certificates, a method of further enhancing the small size of ECC signatures and reducing the overall size of, and the overhead required, to handle the digital certificates that must be transmitted and stored by vehicles and roadside equipment. This feature is used extensively to support privacy by issuing thousands of unique certificates to every vehicle. The SCMS design supports a distributed architecture which allows separation of roles so that automobile manufacturers (OEMs), fleet owners, and Intelligent Transport System (ITS) managers can operate independently while sharing a common root of trust. The goal of the system is to ensure that all vehicles and roadside equipment can reliably exchange authenticated messages even when they are managed by different organizations which do not have existing formal relationships. The SCMS was originally presented by Whyte et al. [1], and was further developed by the Crash Avoidance Metrics Partnership (CAMP). The specifications published by CAMP have since been adopted by the IEEE 1609 Working Group for standardization, with an initial draft planned as IEEE 1609.2.1.

Figure 1 demonstrates the proposed entities within the SCMS and maps the typical communication patterns between them. The roles of each entity are briefly described below.

Policy & Technical Management

In the SCMS, these components manage the overall system. Policy updates are published in a standard format, certified by the Policy Generator, and made available to all participants. High-level controlling entities called Electors are established in order to facilitate Root Management.

Root Management

Establishes the system-wide root of trust, the Root CA. The Electors endorse the Root CA and provide a mechanism for adding or removing Root CAs from the system. The system can have a single Root CA or multiple Root CAs.

Local Management

The Root CA is kept offline and is typically only brought online to certify Intermediate CAs (ICAs), which in turn certify the other components in the system.

Enrollment

The enrollment process takes place before the end-entity starts operation. It involves initial bootstrapping of the end-entity at the OEM, where an out-of-band communication channel is established between the end-entity and the Device Configuration Manager (DCM) in a secure environment during manufacture. The DCM acts as an interface between the end-entity and the Enrollment CA (ECA) and ensures that only authorized devices are able to submit an enrollment request. The enrollment certificate serves as an end-entity's "entry ticket" to request pseudonym certificates. Enrollment certificates can be blacklisted to prevent the request of pseudonym certificates at a Registration Authority (RA).

Pseudonym Issuing

The process of generating pseudonym certificates for the end-entities (EEs) is distributed across four components. This is done so that no single entity knows which pseudonym certificates end up with a specific vehicle, hence; SCMS provides privacy protection against insider as well as outsider attacks. The entities involved in generating pseudonym certificates are as follows.

- Registration Authority (RA): Receives the certificate requests from EEs and forwards them to the Pseudonym CA after expanding the individual request per EE into thousands of requests and shuffling them across multiple EEs. These requests are combined with inputs from Linkage Authorities. The RA knows the enrollment information associated with an EE, and can reconstruct batches of certificates, but cannot see the individual certificates because they are encrypted directly to the EE.

- Pseudonym CA (PCA): The actual issuing CA which collects the information to be inserted in the certificates, signs and encrypts them without being able to trace them to a particular EE. The PCA sees all pseudonym certificates, but does not know which belong to a given batch, or which vehicle is requesting them.
- Linkage Authorities (LA): These components generate linkage values to be included in the EE certificates and sends these linkage values encrypted to the PCA. The LAs do not learn which certificates contain which linkage values. The origin (seed) of these linkage values is only revealed in the case of misbehavior of an EE; this is done by including the seed in a Certificate Revocation List (CRL). Both LAs must cooperate in order to link certificates together. Revealing the seed linkage value, revokes all the pseudonym certificates of the misbehaving EE.
- Location Obscurer Proxy (LOP): As a final privacy-protecting measure, an EE's physical location is obscured to the RA by connecting via a proxy.

Misbehavior Detection

Receives reports of anomalous behavior, identifies bad actors, cooperates with the Linkage Authorities, and creates and distributes CRLs. If both LAs reveal the linkage seeds used to generate a batch of certificates (only at the request of the Misbehavior Authority), then they can be published in a CRL, thereby revoking the entire batch of pseudonyms.

End-Entity

An End-Entity (EE) may be On-Board Equipment (OBE) on the vehicle or Roadside Equipment (RSE). In general, OBEs are issued batches of pseudonyms so that they can rotate through them for privacy reasons, while RSEs are issued individual short-lived single-purpose pseudonyms, as privacy is generally not necessary for RSEs.

3 Overview of Misbehavior Detection

The SCMS is purpose-built to facilitate message authentication and authorization without revealing the identity of users, or allowing the activity of a single user to be tracked over time. There are two fundamental assumptions accounted for in the SCMS design:

1. Devices should continue to function even if they cannot communicate with the SCMS for a long period of time (up to 3 years)
2. No single component should be able to compromise individual privacy in any way

This is achieved by issuing devices 20 pseudonyms per week for 3 years into the future, with the pseudonym issuing process split across five different entities as described in the previous

Section. These assumptions can be restated from the perspective of misbehavior management:

1. If a misbehaving device is not actively revoked from the system, it can continue to misbehave for up to 3 years
2. No single component (such as the Misbehavior Authority) has enough information to add a device to a revocation list (CRL) on its own

Note that the first assumption can be considered somewhat contradictory, as vehicles are assumed to be offline for long periods of time, but misbehavior detection, reporting, and response require that vehicles be online to forward misbehavior reports and to receive misbehavior-related updates.

Misbehavior detection and response consists of a sequence of steps:

1. Detection at the device level (vehicles or infrastructure)
2. Reporting to a central authority (the MA)
3. Analysis and correlation of reports (by the MA)
4. Enacting a decision
   I. Requesting linkage information from the LAs
   II. Blacklisting the device at the RA (invalidate enrollment info)
   III. Revoking the device (invalidating its downloaded pseudonyms)

Each of these phases is described in detail below.

3.1 Detection

There is no complete list of behaviors that should be monitored and reported as misbehavior. Some potential first steps for what might constitute misbehavior within safety applications using the SCMS could be:

- Device using expired credentials
- Device using incorrect credentials (e.g. bad permission, geo-fence)
- Device sending messages with incorrect signatures
- Device sending messages with invalid location
- Device sending messages with other invalid data (e.g., time)

These all represent situations where a device is consistently broadcasting information that does not pass standard validation, or that is not consistent with the sensor readings of the receiver. Because these situations can all be detected and reported by a single observer, they are not difficult to implement. Similarly, reports from multiple observers can be combined to provide assurance that the reports are accurate. Combining data from multiple sources is essential, as

end-entities have multiple concurrently valid pseudonyms they could use to falsely sign misbehavior reports.

Some initial research has been performed at the University of British Columbia on the problem of "spectrum interference", wherein a device maliciously uses the wireless communication channel in a non-standard way in an attempt to degrade or prevent use of the channel. Misbehavior of this sort could be detected at the device level and reported for analysis.

3.2 Reporting

Once sufficient evidence of misbehavior has been collected by a device, it is formatted into a misbehavior report and forwarded to the MA, via the RA, for analysis. At the time of writing, there current specification for misbehavior reporting are provided as ASN.1 modules as part of the CAMP protocols [2] used to support connected vehicle pilot test sites in the United States. The most recent update to the ASN.1 defined misbehavior reports as follows:

```
MisbehaviorReportContents ::= SEQUENCE {
    version             Uint8(1),
    generationTime      Time32,
    policyFilename      PolicyFilename,
    reportType          ReportType,
    evidentiaryData     SEQUENCE (SIZE(1..3)) OF Evidence,
    ...
}
```

This structure defines the report as containing a version number, a timestamp, the current policy file in use by the device, and the core misbehavior data: The type of misbehavior, and the evidence to support this claim.

```
ReportType ::= CHOICE {
  proximityPlausibility      ProximityPlausibility,
  warningReport              WarningReport,
  ...
}
```

The type of misbehavior in the current specifications are limited, with "proximity plausibility" being the primary example. This is likely to encapsulate location-related misbehavior observations.

```
Evidence ::= SEQUENCE {
  observedNeighborList  SEQUENCE (SIZE(0..MAX)) OF
SignedBSMsWithCertificate,
  reporterBSMs          SEQUENCE (SIZE(1..10)) OF SignedBSM,
```

```
  suspectVehicleList    SEQUENCE (SIZE(1..10)) OF
SignedBSMsWithCertificate,

  ...

}
```

The evidence supporting a report consists of three fields. The reporter includes a set of signed BSMs that they were transmitting at the time of observation. This provides a baseline for comparison against the claimed misbehaving messages, which form the second piece of evidence. Thirdly, the vehicle includes BSMs from other nearby vehicles to corroborate its obserbations and to prove the reported vehicle is the one with anomalous data. To conserve bandwidth, the full pseudonym certificate is only included periodically in BSMs. For misbehavior reporing, messages that contain full pseudonym data are necessary.

Note that the messags presented above are not standardized outside of the ASN.1 definitions given by CAMP. The above definitions were used to support CV-Pilot sites in the United States, but a slightly different version of the ASN.1 has seen wider adoption outside of the CV-Pilots. Standardization of the CAMP specifications is currently underway by IEEE, to be initially published as IEEE 1609.2.1, and additional changes are likely to be made during this process.

3.3 Analysis

The actions performed by the MA in order to determine whether a set of reports constitutes misbehavior or not is not currently defined and is an active area of research. Fundamentally, the MA has a few tools at its disposal. First, it can correlate reports from multiple senders in similar geographic regions to determine if reports are likely to belong to the same vehicle. Alternatively, it can search for distinctive features in reported messages to potentially correlate them. Once the MA has grouped misbehavior reports into groups likely to belong to the same vehicle, it can reach out to the PCA to determine which contributions from each LA were combined in each suspected pseudonym, and then to the LAs individually to determine if a set of values belong to the same vehicle or not.

Once correlation information is received from the LAs, the MA can either refine its queries in response, or make a decision as to whether the reported messages constitute misbehavior.

3.4 Blacklisting and Revocation

If the MA determines that misbehavior has occurred and revocation is in order, then blacklisting and revocation processes are carried out.

Blacklisting refers to the invalidation of a device's enrollment certificate. This is the credential used to authenticate with the RA when requesting and downloading pseudonyms. Armed with the linkage information from the LAs, the MA provides the RA with sufficient information for it

to determine which enrollment certificate was associated with the request for the misbehaving pseudonyms, and instructs the RA to blacklist this certificate. Once blacklisted, the RA will no longer offer services to the blacklisted device. This prevents the download of any previously requested pseudonyms, or requests for new pseudonyms. The blacklisted vehicle is effectively unable to communicate with SCMS components at this point.

Blacklisting cannot prevent a vehicle from using pseudonyms it has already downloaded. In order to inform other participants in the network that a set of pseudonyms have been revoked, the CRL must be updated with the linkage information associated with the revoked vehicle. Linkage information starting with the current week is published on the CRL for distribution to all participants in the SCMS. Linkage information is computed using a "forward only" process, which means it is possible to publish a value on the CRL that enables all current and future pseudonyms of a revoked vehicle to be identified and linked, but does not retroactively compromise privacy for actions prior to the time of revocation.

4 Sybil Attacks

A Sybil attack [3] refers to an attack where a single user of a system presents multiple distinct identities in order to amplify its capabilities or evade detection. By design, every vehicle using the SCMS has 20 concurrently valid pseudonyms which are not linkable by any user or component of the system. Therefore, a malicious user can authenticate using 20 different identities under normal circumstances, and no single entity can distinguish these from 20 distinct vehicles. This is one of the reasons the MA is empowered to order the LAs to reveal if two or more provided pseudonyms are linked to the same vehicle or not when analyzing misbehavior reports.

Multiple identities are not just an issue when it comes to committing misbehaving acts. A misbehaving vehicle can not only pretend to be 20 different vehicles when sending misbehaving messages, but can also submit misbehavior reports framing an innocent vehicle using its 20 pseudonyms. Messages from an honest vehicle can be falsely reported as misbehavior even if it is valid. A malicious vehicle simply needs to behave as though it is 20 different vehicles in another location and submit 20 misbehavior reports claiming the honest vehicle is reporting incorrect locations. This observation suggests that the MA must investigate both the reporter and the reported vehicles when misbehavior is reported. Therefore, the act of reporting misbehavior requires an honest user to weaken their privacy within the system. This could incentive privacy-conscious users to avoid submitting misbehavior reports, potentially weakening the integrity of the system.

A potential workaround to this may be to leverage connected infrastructure as the primary source of misbehavior reports. Non-private participants of the system, such as traffic lights, are likely to be more closely monitored, possess only a single valid certificate at a time, and do not have privacy concerns that could prevent reporting.

5 Issues with Scalability

When a vehicle is has its credentials revoked it is placed on a CRL, which is periodically distributed to participants of the network. Because of the fundamental assumption that vehicles may be offline for up to 3 years, distribution of this CRL cannot always be guaranteed in a timely manner. Updates can be sent in a peer-to-peer manner, but it is currently unknown how well such an approach scales and how quickly updates can be propagated through the system.

Each entry on the CRL represents a "seed" value that was used to compute the linkage information for all 20 pseudonyms in that weekly batch. To keep the CRL as short as possible, only the seed is published, and each receiver expands the seed into the 20 linkage values that can be used to link and identify pseudonyms for that week. The resulting computed CRL is necessarily 20 times larger than the distributed CRL. In a full-scale deployment, vehicles will be receiving hundreds of BSMs per second, each of which should be checked against the CRL before being considered for safety purposes. If more than a handful of vehicles are revoked, then the CRL could grow to hundreds, or even thousands of entries very quickly, which is likely to have performance impacts on the verification of messages. An adversary could attack the system by attempting to cause a large number of revocations

Some computational optimizations and technical workarounds are possible to speed up processing of the CRL, but ultimately the only protection against attacks on scalability is judicious application of the revocation process. Alternatively, the assumption that vehicles are offline for years at a time could be revisited. If vehicles are able to update their pseudonyms more frequently, then pseudonyms do not need to be issued far into the future and the system may be able to pivot to passive revocation, wherein pseudonym are not revoked, but a misbehaving vehicle is prevented from obtaining new pseudonyms in the future

6 Future Considerations

Connected vehicle applications are still in an early phase of development. Section 3.1 considered some early candidates for misbehavior detection, but as connected vehicle

technology is more widely deployed and additional applications are built on the platform, new behaviors and definitions of misbehavior are certain to emerge. It follows that misbehavior definitions in vehicles should be implemented in a manner where updates are possible.

The reporting and analysis of misbehavior are still in a very early stage of development. While the SCMS is being standardized, early efforts are focused in the same areas that CAMP focused on, namely vehicle enrollment and provisioning. Misbehavior is unlikely to see substantial overhaul or update until later versions of the standard. While this is limiting in the short term, it suggests that we are in a critical period where large updates to standards are still possible. New notions of misbehavior, such as spectrum misbehavior, can be captured before misbehavior standards solidify.

On top of the standards gaps in the area, misbehavior is a sensitive process that can potentially weaken the entire SCMS if not handled correctly. This paper considered Sybil attacks and the implications of the SCMS's connectivity model on privacy, as well as scalability issues with the revocation process that could prove fatal against a determined attacker. Policies around what processes and thresholds the MA uses to investigate misbehavior will need to be carefully designed in a manner that ensures the SCMS remains robust against attack.

7 References

[1] William Whyte, André Weimerskirch, Virendra Kumar, and Thorsten Hehn, "A security credential management system for V2V communications", IEEE Vehicular Networking Conference, 2013.

[2] CAMP LLC, "Crash Avoidance Metrics Partnership (CAMP) CV-Pilots Documentation", https://wiki.campllc.org/display/SCP

[3] John Douceur, "The Sybil Attack", International workshop on peer-to-peer systems, Springer, Berlin, Heidelberg, 2002.