Efficient Mechanisms to Provide Convoy Member and Vehicle Sequence Authentication in VANETs

Ahren Studer, Mark Luk, Adrian Perrig
Electrical and Computer Engineering
Carnegie Mellon University
{astuder, mluk, adrian}@ece.cmu.edu

***Abstract*— Vehicular Ad hoc Networks (VANETs) are on the verge of deployment. In the near future, wireless vehicle-to-vehicle and vehicle-to-infrastructure communication will enable numerous safety, convenience, and business applications. Security is a necessary pre-requisite for adoption of these technologies.**

As we demonstrate in this paper, VANETs require two new security properties: Convoy Member Authentication (CMA) and Vehicle Sequence Authentication (VSA). These security properties detect a range of VANET attacks. We propose novel protocols that provide CMA and VSA. We analyze and evaluate our protocols and conclude that they represent an important step towards enhancing VANET security.

I. Introduction

In 2005, there were 43,443 traffic fatalities in the United States alone [17]! The government and many manufacturers and are pushing for increased safety mechanisms in vehicles to address the rising number of fatalities and to reduce the $260 billion spent annually on accident-related healthcare [6]. Current state of the art automotive safety solutions use range finding lasers and other expensive hardware to provide drivers of high-end vehicles with more information about their surroundings. Within five years, manufacturers will deploy vehicles with dedicated short range communication (DSRC) capabilities at a fraction of the cost of today's safety solutions to provide the same functionality. DSRC allows a vehicle's On Board Unit (OBU) to communicate with other OBUs and Road Side Units (RSUs) to form a Vehicular Ad Hoc Network (VANET). In addition to safety applications, VANETs will provide convenience and commercial applications to reduce time on the road and to improve driving experience. Given the highly safety-sensitive nature of VANETs and the risks associated with their wireless communication, it is clear that we need to secure these networks against adversarial activity.

Manufacturers will deploy a number of safety applications once VANETs become available [1]. These safety applications include: Electronic Emergency Brake Light (EEBL), Road Hazard Condition Notification (RHCN), Road Feature Notification (RFN), Slow Vehicle Alert (SVA), and Post Crash Notification (PCN). These applications help alert other drivers of dangerous situations or conditions. EEBL alerts drivers when a vehicle rapidly decelerates, to reduce the chance of rear-end collisions. RHCN broadcasts alerts about debris (e.g., ice or trash) on the road. RFN alerts drivers when they approach a steep hill or a section with a notably lower speed limit (e.g., a school). SVA and PCN alert drivers of a slow vehicle or a possible crash in the lanes ahead. Alerts from these VANET applications provide drivers more time to react to dangerous conditions, reducing the chance of an accident.

Our key insight is that VANET safety applications require authentication of the physical properties of the sender for security, not just traditional cryptography based identity authentication. Safety alerts are only relevant if the braking car, patch of ice, or accident is on the road in front of the recipient in what we call the Area of Relevance (AOR). A malicious entity could falsely claim a position in front of recipients and broadcast fake alerts as a way to disrupt traffic. To identify cars in the AOR, we introduce Convoy Member Authentication (CMA) and Vehicle Sequence Authentication (VSA) to verify the sender of the alert is traveling with and in front of the recipient. If OBUs have CMA and VSA, attackers can only fool a recipient with a fake alert while physically in a victim's AOR.

Although several VANET security mechanisms have been proposed [10], [11], [18], [20], [21], [27], none of the proposed mechanisms that we are aware of address all of the requirements needed to secure VANET safety applications. These works focus on authentication of the identity of another OBU, rather than authentication of the validity of the alert. More concretely, under previous approaches an attacker could easily sign a spurious safety message, which could cause drivers to unnecessarily apply their brakes or change lanes.

In essence, the problem we address in this paper is to verify that a given message indeed originates from a legitimate vehicle driving on the same road ahead of the recipient's current location (i.e., inside the Area Of Relevance (AOR)). Simultaneously, this security mechanism will also provide a useful filter against non-malicious useless messages: e.g., a braking alert from a vehicle driving on a nearby road or in the opposite direction on the same road. Since DSRC messages have an expected range on the order of 300 meters, such a filter is critical to avoid spurious false alarms. More specifically, there are three attacks in particular that we will defend against in this paper: an attacker that attempts to inject alerts in opposing traffic, a stationary attacker on the side of the road that tries to

inject bogus alerts, and an attacker driving on the road who is trying to fake an alert to vehicles ahead of it, claiming to drive in front of them.

Unfortunately, current approaches for secure positioning [4], [12]–[15] address the dual problem, where a node attempts to correctly determine its own location despite the presence of an adversary, as compared to the challenge in VANETs where nodes know their own location and wish to verify the location of other nodes. Mechanisms for secure location verification [2], [3], [23] represent a more general approach for location claim verification, however, prior approaches make use of trusted infrastructure which may not be available in a VANET context. Moreover, the location verification problem we address in this paper is a significantly simplified problem because we only need to verify whether the vehicle is in front or behind with respect to a line perpendicular to the current direction – thus, we hope to achieve a much more efficient mechanism.

The IEEE 1609.2 standard [11]. The current IEEE standard provides guidelines for secure message formats and how to process those messages in VANETs. This information is an important step when designing systems to operate in environments with entities from several manufacturers. However, the IEEE standard does not provide any specific protocol. The general framework suggests the use of a Public Key Infrastructure. Unfortunately, a PKI does not fulfill the security requirements of many VANET safety applications because *identification* of the vehicle that sends a message through authentication or signature techniques is often unnecessary in VANETs. Instead, for safety messages, the important property that needs to be verified is that the sender is a legitimate vehicle driving on the same road ahead of the receiver. The current standard implies that digital signatures provide all of the authentication needed for vehicle-to-vehicle and vehicle-to-infrastructure (i.e., RSUs) communication. Unfortunately, the current standard provides no provisions for verifying that a sender is a legitimate vehicle driving on the road. Similarly, for safety applications, the identity of the signer is often not the important property to authenticate, but location and movement of the signer needs to be verified. Mechanisms that verify these physical properties will filter spurious alerts and detect a variety of malicious activities. With such mechanisms in place, attackers are forced to drive in the AOR of the vehicles they desire to attack.

Thus, the focus of prior work on *vehicle identification* is insufficient and not even necessary for several applications.

Paper Contributions. The main contributions of this work include:

- We observe that identification of vehicles is of little importance to VANET safety applications. Instead, we propose that the physical location and movement of the sender requires verification.
- We provide formal definitions for the physical properties needed to help secure VANET safety applications. Convoy Member Authentication (CMA) allows a vehicle to determine which vehicles are driving in the same direction on the same road. To determine the order of vehicles on a road, we introduce Vehicle Sequence Authentication (VSA).
- We present BCMA (Beacon-based Convoy Member Authentication), a mechanism to provide CMA and TVSA (Timing-Based Vehicle Sequence Authentication), novel security mechanisms to provide VSA.
- We analyze, implement, and evaluate our security mechanisms in a realistic VANET simulator.

Outline. In Section II, we provide formal definitions for our attacker model, CMA, and VSA, and state our assumptions. We introduce our mechanism to provide CMA in Section III. Section IV introduces our mechanisms to provide VSA. We present simulation results of our CMA and VSA mechanisms in Section V. We discuss how our work relates to previous publications in Section VI, and make concluding remarks in Section VII.

II. Problem Definition

In this section, we provide a concise problem definition, state our assumptions, and present our attacker model.

A. Problem Definition

Numerous papers have discussed the high level security requirements of VANETs such as authentication of parties, or how to provide privacy. In this work we focus on the physical properties needed to ensure proper operation of safety applications. More specifically, the challenge is to verify that the source of an alert is driving on the same road and in the same direction as the recipient (Convoy Member Authentication (CMA)), and that the sender is ahead on the road (Vehicle Sequence Authentication (VSA)). Figure 1 shows how the combination of CMA and VSA helps identify which vehicles are in a region we call the Area of Relevance (AOR). The size of the AOR should change with speed to reflect how long it will take a vehicle to reach the site of an alert. For example, a vehicle on the highway traveling at 110km/h will have an AOR that includes all of the lanes of traffic heading the same direction 300 meters in front of the OBU (or roughly the region the vehicle may traverse in the next 10 seconds). The same vehicle in an urban environment may slow down to 40 km/h. At this slower speed, the AOR will decrease to only include any traffic traveling the same direction as the OBU and within 110 meters ahead of the vehicle. If the OBU provides the driver with numerous spurious alerts from outside the AOR (e.g., debris on the other side of the road, a braking vehicle a significant distance ahead on the road, or a crash behind the vehicle), the driver will start ignoring the OBU and the applications will fail to improve roadway safety. Instead, we must provide mechanisms that authenticate physical properties in the face of inaccurate location claims from malfunctioning OBUs or malicious parties trying to cause confusion.

To help OBUs identify which vehicles are traveling together on the same road we propose the Convoy Member property. We formally define a group of vehicles traveling together in the same direction on the same road as Convoy(α, β). The same

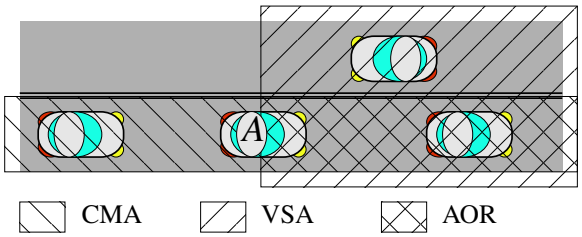

Fig. 1. How vehicle *A* combines CMA & VSA to form the AOR.

road is defined as any lanes of traffic without a physical barrier between them. If OBUs are sparse, the convoy consists of every vehicle within a radius of α meters traveling the same direction on the same road. The value of α changes with vehicle speed to reflect the area the vehicle may encounter in the next ten seconds. If OBUs are dense, the OBU only monitors the β closest vehicles traveling the same direction on the same road. We can afford to only monitor the β closest OBUs because safety messages do not need to propagate as far when traffic is congested and moving slowly.

In Section III, we introduce our mechanism to provide Convoy Member Authentication (CMA). A *positive* for CMA is the detection of an OBU that falsely claims a position in the convoy. An accurate CMA mechanism has a low probability of excluding legitimate vehicles from the convoy (low false positive), and a high probability of detecting vehicles that incorrectly claim to be part of the convoy (high true positive).

CMA alone does not fulfill the requirements of some safety applications. For example, it sounds reasonable for the OBU to alert the driver whenever a member of the convoy generates an Electronic Emergency Brake Light alert. However, what if the sender of that message is a vehicle traveling well above the speed-limit on the highway and is about to rear-end a vehicle driving below the speed-limit? If the slow vehicle brakes in response to this EEBL message from *behind the vehicle*, the chance of an accident is increased. The safest action an OBU could take may be to simply let the driver continue as though no alert was ever received. This example demonstrates the need for Vehicle Sequence Authentication (VSA).

Vehicle Sequence defines which vehicles are in front of or behind the current OBU. For example, VSA detects and ignores a malicious party behind the OBU that claims a location ahead of the OBU. A *positive* for VSA is the detection of an OBU that makes location claims which contradict the true vehicle sequence. An accurate VSA mechanism has a high probability of detecting an OBU that makes a location claim that breaks vehicle sequence (high true positive) and a low probability of incorrectly labeling a vehicle which makes legitimate location claims (low false positive).

B. Assumptions

In this work we make assumptions about the key management in VANETs, presence of other VANET applications, and the capabilities of the VANET participants.

We assume every OBU possesses an Elliptic Curve Cryptography (ECC) public/private key pair K^{+}_{OBU}/K^{-1}_{OBU} and a certificate from a trusted authority (which has a public key K^{+}_{CA} trusted by all OBUs) to prove the validity of the public key and to tie the vehicle's identity to the public key $\{Id, K^{+}_{OBU}\}_{K^{-1}_{CA}}$ as proposed by the IEEE 1609.2 standard [11]. OBUs will digitally sign each message using the Elliptic Curve Digital Signature Algorithm (ECDSA) so recipients can verify the message was not tampered en route. In addition, we assume the key management system in VANETs will provide timely revocation and verification of OBUs' keys, such that a single vehicle can only have a single valid key at any given time. With only a single key, malicious nodes will only have a single identity and will not have an unfair advantage in protocols that use voting to determine crucial values.

We assume that every T seconds each vehicle will generate a signed beacon that includes the vehicle's position and trajectory. The purpose of these beacons is twofold. One, the VANET safety application Cooperative Collision Warning (CCW) uses these broadcasts to determine when vehicles are about to collide. Two, beacons provide additional information about the (true or claimed) location and trajectory of vehicles on the road. If no beacons were used, a single alert when an event occurs (e.g., a RHCN alert about debris on the road) may not provide receiving vehicles enough information or time to determine if the alert is relevant and legitimate.

We assume that legitimate nodes have correct location information. GPS provides location information within a few meters. However, GPS signals are not authenticated and are thus susceptible to spoofing. One simple mechanism to thwart GPS spoofing is through using map information and dead reckoning.[1] Given a correct initial position estimate, high-resolution map information and local trajectory information, dead reckoning provides a means to estimate the position despite intermediate lack of GPS information and to detect and filter out spoofed GPS information.

C. Attacker Model

To focus our discussion, we consider a specific attacker model against the safety applications discussed in the introduction. From a high level, the basic security requirement is that the message originates from a vehicle in the Area of Relevance (AOR) (thus, all EEBL or RHCN messages originating from vehicles outside of the AOR should be ignored or weighed with less importance). Without a secure mechanism in place, attackers positioned outside of the recipients' AORs could fool drivers with malicious safety messages.

In this paper we deal with three specific attacks where vehicles falsely claim to be in the AOR: an attacker in opposing traffic that claims to be driving with the vehicle, an attacker on the side of the road that claims to be a legitimate vehicle, and

[1]From Wikipedia.org: "Dead reckoning (DR) is the process of estimating a global position of a navigating agent by advancing a known position using direction, speed, time and distance of travel."

an attacker behind the receiver that claims to be in front of the receiver.

We assume attackers have valid ECC keys and certificates, are polynomial-time limited in computation, have limited control over the wireless network, and constitute a small fraction of the population of VANET participants. An attacker's valid credentials allow it to generate and sign VANET messages such that recipients can verify the signature through the VANET PKI. The attackers are polynomial time bounded so hard problems that form the basis of public key cryptography (i.e., discrete log) cannot be broken. This prevents attackers from forging signatures for other OBUs or RSUs. Attackers can jam radio signals or use directional antennae to broadcast messages to a subset of the surrounding vehicles. However, the confining nature of roads prevents an attacker with a directional antenna from sending a message to an OBU several cars ahead without allowing closer OBUs to receive the message (i.e., if OBUs *A B C D E* are driving in that order, OBU *B* cannot send a packet to *E* without *C* and *D* hearing it). Finally, we assume malicious parties represent a small fraction of the entire VANET population. If the majority of vehicles were dishonest, law enforcement mechanisms would be effective in curbing malicious behavior.

III. Beacon-based Convoy Member Authentication (BCMA)

Convoy Member Authentication (CMA) allows OBUs to determine what other OBUs are traveling in the same direction on the same road. If VANETs lack a mechanism to provide CMA, OBUs could incorrectly alert drivers when OBUs driving in the opposite direction or radios on the side of the road generate alerts. Vehicles traveling on roads have highly constrained trajectories; other vehicles traveling on the same road in the same direction are often in close proximity for extended periods of time. Thus, if we continuously receive beacons from other OBUs during a time period, we believe they are driving in the same direction.

Exploiting this observation, we propose the Beacon-based Convoy Member Authentication (BCMA) protocol. BCMA relies on continued presence to determine if a vehicle is indeed in the vicinity for an extended time period. Continued presence is defined through the use of a required number of beacons before a vehicle is accepted as part of the convoy (i.e., a vehicle traveling with the message recipient). In BCMA, a vehicle only considers another vehicle part of the convoy after it hears a threshold τ beacons during $T \cdot (\tau + x)$ seconds, where T is the minimum time between CCW beacons and x is the maximum number of acceptable lost beacons. Note that beacons from a single OBU that are more frequent than T are ignored. The assumption here is that a vehicle traveling in the opposite direction or a stationary radio will be in radio range for a shorter period of time when compared to vehicles traveling in the same direction. This idea is similar to the work by Golle et al. [8] where OBUs build a model of the VANET environment and select the most probable (e.g., nodes traveling together will hear more of each others beacons over a period of time). With BCMA in place, vehicle *A* in Figure 4 will recognize *M* as not belonging to the convoy, since *A* will not hear τ or more beacons during the time when *M* enters and leaves *A*'s radio range. The value of τ depends mainly on the vehicle's speed. With a large τ, vehicle *D* may ignore *B*'s alerts if *B* recently merged onto the road (a false positive). With a small τ, the OBU may incorrectly believe slow oncoming traffic or a radio on the side of the road is part of the convoy (a false negative). This mechanism only provides CMA, and thus vehicle *B* in Fig. 4 would believe *M* is in *B*'s AOR if *M* claimed a position in front of *B*. However, this is a failure of vehicle sequence authentication not CMA.

BCMA is meant to achieve convoy member authentication and thus determine which vehicles are traveling in the same direction. The assumption here is that only convoy members remain in a recipient's radio range for an extended period of time and can reach the recipient with broadcast beacons for τ successive intervals.

A. BCMA Security Analysis

BCMA is a crude heuristic to efficiently filter out the majority of malicious locations claims. However, attackers may still reach an OBU with the necessary τ beacons without driving in the convoy. For example, an attacker traveling in the opposite direction or stopped on the side of the road could use a large transmission power to try and defeat BCMA. An attacker could also travel on a parallel road in the same direction and pass BCMA's convoy membership test. VANETs should deploy TVSA in addition to BCMA for higher resilience to such attacks.

IV. Timing-Based Vehicle Sequence Authentication (TVSA)

Convoy Member Authentication allows OBUs to determine which vehicles are traveling in the same direction on the same road. However, when an OBU receives a safety message, the OBU needs more specific information to determine if the sender is in the AOR, in particular, which vehicles are in front of or behind the OBU. For the safety applications discussed in the Introduction, only alerts from vehicles ahead are of use; debris or a crash behind the vehicle is irrelevant. In a benign environment, alerts could include the OBU's current location and velocity, which would suffice for determining whether the sender is in front or behind. However, an attacker could claim a position further ahead, generate a false alert, cause an accident, and try to collect insurance money or sue for more. We propose Timing-Based Vehicle Sequence Authentication (TVSA) as a simple, yet powerful mechanism for Vehicle Sequence Authentication (VSA). As the name implies, TVSA uses beacon reception time information to determine the true sequence of vehicles on the road. We recognize that time-of-flight is a popular mechanism to verify location and review related work in Section VI. First we present Timing-Based Vehicle Sequence Authentication-Global Synchronization (TVSA-GS), where OBUs use nanosecond time synchronization to

verify VSA. Perfect synchronization is difficult to achieve, but GPS can provide time information within ±20ns [16]. In Section IV-B, we present Timing-Based Vehicle Sequence Authentication-No Synchronization (TVSA-NS) which utilizes other OBUs to determine the difference between two OBUs' internal clocks when global synchronization is unavailable. We include an analysis of the mathematics that allow both types of TVSA to detect vehicles that try to break vehicle sequence. Section IV-D performs a security analysis of TVSA-GS and TVSA-NS.

A. TVSA-Global Synchronization (GS)

TVSA uses physical limitations and an honest majority to determine the sequence of vehicles on the road. The intuition behind TVSA is that the difference between beacon arrival times at different locations on the road can reveal the true location of the source of the beacon. An overview of TVSA is as follows. All vehicles provide a location claim in their periodic beacons. To verify a location claim, the recipient acts as the "verifying vehicle" and uses third parties (e.g., another OBU or RSU) to acquire additional witness data. The verifying vehicle estimates when the vehicle in question broadcast the beacon based on the reception time of the beacon and the distance from the claimed location. Next, the verifying vehicle compares this broadcast time estimate with other vehicles' estimates of the original sender's broadcast time. If the estimates disagree by more than a threshold amount, the original sender must have lied about its location claim and violated the true vehicle sequence. Instead of having the verifying OBU query other OBUs for witness values, every OBU includes timing and distance information in every beacon message in order to act as a witness to every other OBU. More specifically, OBUs periodically broadcast beacons with several pieces of information: location and velocity, local arrival time of other beacons, and the relative distance from the beacon sender's claimed location when the beacon was received. Each OBU that receives a beacon will record the local arrival time ($t_{SenderReceiver}$) and how far away it was from the claimed location ($Dist(Sender, Receiver) = Dist(Receiver, Sender)$), and broadcasts that information in its next beacon. MAC layer timestamping [7] can provide the level of accuracy necessary when OBUs record reception times. Without MAC layer timestamping, network stack reception delays will vary greatly across vehicles and make TVSA ineffective.

For example, in the scenario in Figures 2, 3, and 4, B acts as the verifying vehicle for M's location claim. To verify M's location claim, B checks that M's beacon arrived at an appropriate time at each witness given each witness' location, M's claimed location, and the times of reception. To check location claims, in Step ☆ in Figure 3, B verifies that B's assumption for M's broadcast time ($t_{MB} - \frac{Dist(M,B)}{c}$, where c is the speed of light) matches witness W's assumed broadcast time ($t_{MW} - \frac{Dist(M,W)}{c}$) plus or minus some acceptable difference Δ. If every OBU's internal clock is synchronized and no vehicle makes false location claims, the two values will be equal. The level of OBU synchronization possible dictates the value of Δ. If tight time synchronization is possible, a small Δ will allow TVSA to detect false sequence claims without mislabeling legitimate claims as malicious. If synchronization between OBUs is loose, Δ must be set large to prevent mislabeling of legitimate vehicles (false positives). However, if TVSA with a large Δ analyzes the claims, verifying vehicles will not be able to detect malicious parties that claim a false vehicle ordering (false negatives).

Note: For clarity, authentication and other data in the packet have been excluded.

$M \rightarrow *$: (M, Loc_M, Vel_M)
M broadcasts its location and velocity.
$B \rightarrow *$: $(B, Loc_B, Vel_B)\{(M, Dist(M,B), t_{MB})\}$
B broadcasts its location, velocity, when it heard M's beacon, and the relative distance at that time.
$C \rightarrow *$: $(C, Loc_C, Vel_C)\{(B, Dist(B,C), t_{BC}), (M, Dist(M,C), t_{MC})\}$
$E \rightarrow *$: $(E, Loc_E, Vel_E)\{(B, Dist(B,E), t_{BE}), (C...), (M...)\}$
.
: (Vehicles continue to broadcast their beacons and witness values.)

$M \rightarrow *$: $(M, Loc_{M2}, Vel_{M2})\{(B, Dist(B,M), t_{BM}), (C, Dist(C,M), t_{CM}, ...\}$
: Before recording M's new info, each vehicle checks M's old
: location claim.
$*$: $VerifyClaim(M)$ See Fig.3

Fig. 2. TVSA messages in example run

```
    //For each witness B heard
    for(W ∈ VehiclesHeard && W! = M)
      //Compare the estimated broadcast time for M.
☆     if(||t_MB − Dist(M,B)/c − (t_MW − Dist(M,W)/c)|| ≤ Δ)
        VotesForM++
      else
        VotesAgainstM++

    if(VotesForM > VotesAgainstM)
      Trust M's Claim
```

Fig. 3. Code OBU B uses to verify M's claim in TVSA

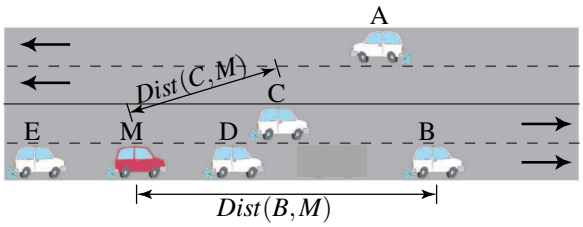

Fig. 4. Example traffic scenario.

Provided tight time synchronization, TVSA uses the distance between OBUs and beacon reception time to detect false vehicle ordering claims. In the next subsection, we discuss how TVSA can work in the absence of time synchronization between OBUs. We conclude with two subsections which discuss exactly how TVSA detects vehicle sequence violations and why a malicious party cannot claim a false vehicle sequence without being detected.

B. TVSA-No Synchronization (NS)

TVSA-GS relies on some source for global synchronization across all OBUs. GPS can provide synchronization or RSUs could act as sources of reference broadcasts [7]. However, physical obstructions can block GPS (e.g., buildings, tunnels, mountains, etc.) and the installation of RSUs on every road would drastically increase the cost of VANET deployment. Dead reckoning can help OBUs maintain accurate information about their own location, but over time, clock drift will affect clock synchronization between OBUs. To address the degradation of clock synchronization, we propose using other witness vehicles' beacons to calculate the clock offset between OBUs.

When V_1 verifies X's location claim in TVSA, V_1 considers all of the different witness reception times and the relative distance between X and the witness at the time of reception. In the absence of time synchronization, V_1 must discover the synchronization error between itself and a witness V_2 ($\delta_{V_1V_2}$). In TVSA-NS, V_1 uses data from another witness which we call a reference vehicle (V_R) to produce $\hat{\delta}^{V_R}_{V_1V_2}$ (the time synchronization error estimate with respect to V_R). V_1 estimates the offset between its clock and V_2's clock using the following equation:

$$\hat{\delta}^{V_R}_{V_1V_2} = t_{V_RV_1} - t_{V_RV_2} - \frac{Dist(V_R,V_1)}{c} + \frac{Dist(V_R,V_2)}{c} \quad (1)$$

The major problem with using other vehicles to provide reference broadcasts is that one (or more) of the reference vehicles may be a malicious vehicle that lies about its location. If we base the error estimate purely on a single malicious party's reference beacon, the estimate will be incorrect. Instead, the verifying vehicle should calculate the synchronization error estimate for each witness with respect to the beacon reception times of *every other witness*. Once the OBU calculates this set of n error estimates, the estimates are ordered from smallest to largest as $\hat{\delta}^1_{V_1V_2}, ..., \hat{\delta}^n_{V_1V_2}$ and the median value ($\hat{\delta}^{n/2}_{V_1V_2}$) is selected as the synchronization offset estimate. Although the mean offset might seem like a reasonable estimate, a malicious reference vehicle's false location claim could lead to a very large or small mean offset. However, the median is only affected after half of the vehicles make false location claims [25]. After V_1 calculates the synchronization offset between itself and the witness V_2 as $\hat{\delta}^{n/2}_{V_1V_2}$, V_1 replaces the inequality in the step denoted "☆" in Fig. 3 with the following inequality to determine if X's location claim is valid.

$$\|t_{XV_1} - \frac{Dist(X,V_1)}{c} - (t_{XV_2} - \frac{Dist(X,V_2)}{c} + \hat{\delta}^{n/2}_{V_1V_2})\| \leq \Delta \quad (2)$$

In this section we explained how other vehicles can help estimate the current witness's synchronization error. If this technique is used, TVSA no longer requires GPS or RSU based synchronization. Provided enough OBUs are present to provide both witness and reference times, TVSA-NS will work everywhere, even in the presence of large time synchronization errors.

C. TVSA Mathematical Analysis

We now discuss how OBUs using TVSA can detect when vehicles violate vehicle sequence. The mathematics presented here are for the one-dimensional case where all vehicles are traveling in a straight line and each car has perfect global synchronization (GS). TVSA operation is the same in the multi-dimensional case where vehicles drive around curves or the road has multiple lanes. However, for simplicity of exposition we limit our discussion here to the one-dimensional case. The mathematics presented here are analogous when global synchronization is not available. However, OBUs require the addition of clock offsets.

Independent of where vehicle X claims to be located, X's true broadcast time (t_X) and when V_1 receives the beacon (t_{XV_1}) are related as follows:

$$t_{XV_1} = t_X + \frac{Dist(X,V_1)}{c} \quad (3)$$

When a vehicle claims location X^*, the recipient's assumed broadcast time (t_{X^*}) changes according to X's location claim:

$$t_{X^*} = t_{XV_1} - \frac{Dist(X^*,V_1)}{c} \quad (4)$$

Combining these two equations, we find the difference between the assumed broadcast time and real broadcast time is

$$t_X - t_{X^*} = \frac{Dist(X^*,V_1)}{c} - \frac{Dist(X,V_1)}{c} \quad (5)$$

If the location claim is on the same side of V_1 as X's real location, (i.e., the vehicle sequence is $X \rightarrow X^* \rightarrow V_1$), V_1's assumed broadcast time for X can be reduced to the following function of X's claimed location (X^*), the true broadcast time, and X's true location:

$$t_{X^*} = t_X + \frac{Dist(X,X^*)}{c} \quad (6)$$

If the location claim is on the opposite side of V_1 and violates vehicle sequence, V_1's assumed broadcast time for X becomes a function of the true broadcast time, the distance between the true (X) and claimed (X^*) locations, and the distance between V_1 and the claimed location.

$$t_{X^*} = t_X + \frac{Dist(X,X^*)}{c} - \frac{2Dist(X^*,V_1)}{c} \quad (7)$$

When V_1 compares reception times and distance claims with V_2, the goal is to determine if the difference between the two assumed broadcast times are within some acceptable range (i.e., $|t_{X_1^*} - t_{X_2^*}| < \Delta$).

TVSA provides more than authentication of the sequence of vehicles on the road. For vehicles within the convoy, the location can be determined. However, we use the term vehicle sequence authentication because vehicles at the start and end of the sequence can make false location claims. A vehicle at the start or end of the sequence can claim a location much further away, or a few meters past the next closest vehicle, and TVSA would not detect the false claim.

The graph on the top of Figure 5 represents the difference between V_1's estimate for X's broadcast time based on X's claimed location and two witnesses V_2 and V_3. As expected, whenever X claims a location on the opposite side of our verifying vehicle V_1, the assumed broadcast times differ. It is important to note that V_1's and V_3's assumed broadcast times agree when the attacker X claims a location behind V_1. Because we use a vote to determine if a location claim is legitimate, V_1 may believe X's location claim behind V_2. However, V_1 will only believe this false claim if V_2 is the only vehicle behind X and there are multiple witnesses in front of V_1. In the same scenario, when V_2 verifies location claims, all of the broadcast estimate differences will show that X's location claims are false, despite the use of voting. The fact that X can claim a location further behind V_1 is acceptable since TVSA is meant to prove or disprove an ordering of vehicles **with respect to the verifier**. As long as V_1 can detect a vehicle whose location claim crosses over V_1 we preserve the true vehicle sequence.

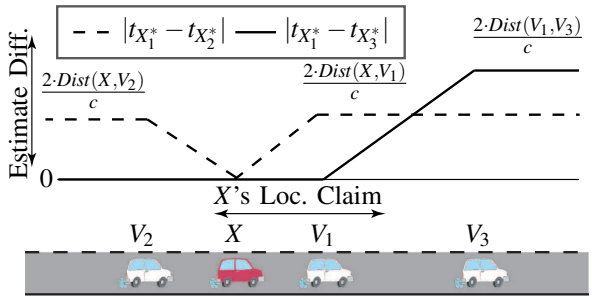

Fig. 5. Broadcast estimate difference vs. varying location claims. ($t_{X_I^*}$ is I's estimate of X's broadcast time)

Even though OBUs never know the true broadcast time and location of a sender, recipients can calculate and compare estimated broadcast times in order to determine if the sender's location claim agrees with or contradicts the true sequence of vehicles on the road. In the next section we discuss how TVSA allows OBUs to determine the true sequence of vehicles on the road, even when a limited number of malicious parties intervene.

D. TVSA Security Analysis

The goal of TVSA is to provide vehicles with a proof of which vehicles are driving in front of or behind them. In a region where OBUs do not purposely lie about their location, TVSA helps isolate malfunctioning OBUs. However, malicious parties may lie about their position. In this section, we examine why an attacker cannot abuse TVSA and claim a false location in a sequence of vehicles or claim a false reception time to slander a third innocent vehicle under the condition that the majority of vehicles are not malicious. However, if the majority of the population is malicious and collude, attackers can violate the true sequence of vehicles.

When an attacker claims a position in a sequence of vehicles, other OBUs use TVSA to check if the vehicle truly is where it claims to be. As discussed in Section IV-C, as long as other vehicles provide accurate location and timing information, such claims will be detected. An attacker controls when its beacon is broadcast, the location claimed in beacons, and the true location of the OBU. Receiving OBUs use their own reception time and distance from the claimed location to estimate the time of broadcast, so TVSA works independent of when an attacker claims a beacon is broadcast. However, an attacker with directional antennas can transmit a beacon at different times in different directions and cause **one vehicle** to believe an incorrect vehicle sequence. For example, if X in Figure 5 wanted to claim a position in front of V_1, X would transmit a beacon forward at t_F and delay a set period of time before transmitting a beacon backwards at time t_R. We assume directional antenna and the nature of roads prevent an attacker from sending a message such that vehicles further ahead on the road receive the message before vehicles directionally in front of the attacker receive the message. In our example, vehicles behind V_1 will agree with V_1's broadcast time estimate for X, but V_3 will have a differing broadcast time estimate (i.e., the solid line in Figure 5 will remain, but the dotted line will become zero). As such, if there are more witnesses behind V_1 than in front of V_1, the attack will succeed. However, the relation between the time of an attacker's transmission forward (t_F) and backward (t_R) depend on the distance between the location of the attacker, its claimed location, and the location of the intended victim in front of it. As a result of this dependency, a t_R that fools one vehicle is different than a t'_R that makes a different vehicle believe the attacker is in front of it. In our example, this limits X to fooling only V_1 or V_3 of a false vehicle sequence. Even in the presence of attackers with directional antennas, TVSA allows the majority of OBUs to detect false location claims that impact the ordering of vehicles.

In TVSA, OBUs rely on other OBUs to provide accurate location and timing information. A malicious party could claim a false reception time as a way to slander a victim and cause a verifying vehicle to doubt the victim's location claim. To mitigate a slander attack, TVSA takes a vote between all broadcast estimate comparisons to determine if a vehicle's location claim is true. Provided there is only a small fraction of the population performing a slander attack, the legitimate votes for a vehicle's location claim should outnumber the malicious parties.

If the majority of OBU's witness values are maliciously fabricated, TVSA will start to validate false location claims that violate vehicle sequence or start discarding true location claims as a result of slander attacks. We assume law enforcement mechanisms would be effective in curbing widespread malicious activities. A single attacker could claim multiple identities to provide multiple witnesses and try to manipulate the voting in TVSA. A Sybil detection mechanism such as [26] can detect the false identities and filter out the invalid witnesses.

Provided the majority of the OBUs in a region are not conspiring, an attacker cannot convince a convoy of OBUs of a location claim that violates the current vehicle sequence or slander another OBU.

V. Evaluation of BCMA & TVSA

We use ns-2 [24] to simulate the different authentication mechanisms from Section III and IV in highway and city settings. Our simulated 1.5 kilometer square 4-lane highway is presented in Figure 6 (a). To represent city traffic we use a traffic scenario generator [22] and the 2 kilometer square city topology presented in Figure 6 (b). In the simulation each OBU has a 250 meter broadcast range and broadcasts two beacons every second ($T = 0.5$). First we describe our simulation environment and the measured quantities. In the following subsections we analyze the different detection capabilities of the mechanisms, how time synchronization impacts TVSA-GS, if TVSA-NS can counteract a lack of synchronization in VANETs, and the overhead associated with our mechanisms.

During simulation we use an Area of Relevance (AOR) that includes every vehicle within radio range traveling in the same direction as the recipient and in front of the vehicle. In simulation, we measure the probability of a legitimate node detecting a malicious entity that claims to be in the AOR when it is not (a true positive) and the probability of a legitimate node ignoring/not believing an alert from a legitimate vehicle in the AOR (a false positive). When analyzing the TVSA-GS as presented in Section IV-A, we assume clocks could be synchronized within $\pm 50ns$, which is a conservative estimate given current GPS system capabilities [16]. This clock error provides some realistic variance in the system that increases the chance of believing attackers' location claims. We simulate three attacks:

(#1) A mobile attacker claims a position with vehicles traveling in the opposite direction.
(#2) A stationary attacker impersonates a vehicle traveling with traffic.
(#3) An attacker traveling on the road claims a position further ahead in the same lane.

The larger dark circles in the topologies (see Figure 6) indicate the locations of attackers' radios for the simulations with a stationary attacker that claims to be moving (Attack 2). Each scenario was allowed to run for 10 minutes of simulated time and repeated several times (5 times for highway simulations and 10 times for city simulations) with the results averaged across all runs to reduce variance. For each simulation, we select a single attacker and a subset of the total nodes at random to generate periodic alerts.

A. General Results

First, we simulated the combination of BCMA & TVSA to determine the detection accuracy of the two mechanisms with varying values for the threshold τ for BCMA. The results of the simulations are presented in Table I. $\tau = 0$ represents TVSA by itself. Excluding attack #1 on the highway and attack #3 in city traffic, TVSA detects over 85% of false location claims. The addition of BCMA helps detect the majority of the remaining false locations at the cost of more false positives.

TVSA has trouble detecting attack #1 in the highway scenario because of oncoming traffic that claims a location ahead of traffic agrees with the true vehicle sequence; the attacker is in front of the verifying vehicle. These errors demonstrate why VSA alone is not enough to properly identify which vehicles are in the AOR; CMA **and** VSA are needed to determine which vehicles are in the AOR. As the threshold τ increases, BCMA identifies which vehicles are part of the convoy, and which are attackers in the opposite direction of traffic. TVSA fails to detect 34% of false location claims in the city because of the synchronization error of $\pm 50ns$. The issue is that vehicles at stop lights or in congestion are close together and the maximum error between broadcast estimates ($\frac{2Dist(V_1,V_2)}{c}$ as shown in Section IV-C) is smaller than the accepted error Δ. Since the error is within the accepted range, the victims believe the location claims that violate the true vehicle sequence.

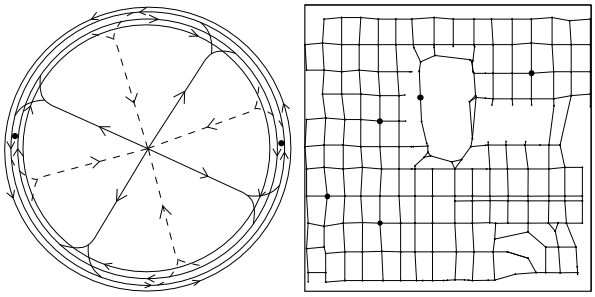

(a) Highway Topology (b) City Topology

Fig. 6. Topologies Used to Simulate Traffic

Scenario	attack	FP/TP Rate		
		$\tau = 0$	$\tau = 2$	$\tau = 5$
Highway Topo.	1	0%/69%	4.5%/82%	10%/100%
40 cars/km	2	0%/90%	6.3%/93%	9.6%/93%
110 km/h (avg.)	3	0%/91%	6.4%/93%	11%/93%
City Topology	1	0%/86%	17%/100%	22%/100%
200 cars	2	0%/95%	16%/99%	22%/100%
55 km/h (avg.)	3	0%/66%	16%/76%	22%/77%

TABLE I
Highway and City Simulation Results

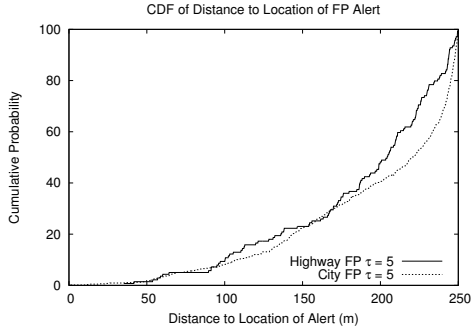

Fig. 7. Distribution of Distance to the Location of False Positives

As expected, the value of τ should be selected based on speed and traffic patterns. On the highway vehicles stay in range for a long time so a τ of 5 improves detection of attack #1 by $\approx 35\%$ with fewer than 10% false positives. The same τ for the city simulation detects 95% attack #1, but causes over 20% false positives. However, the false positive of BCMA are location

dependent. As shown in Fig. 7, only 10% of the false positives correspond to locations within 100 meters of the vehicle. If the AOR were limited to 100 meters in front of the OBU, alerts past 100 meters would be ignored. The probability of a false positive would decrease, and a driver traveling at highways speeds would have at least 3 seconds to respond to the alert (assuming the vehicle's speed is < 120 km/h), an ample amount of time.

Additional simulations were performed with numerous traffic configurations. However, due to space considerations we are unable to include all results. We found that as traffic density increases, network contention causes packet loss. With more packets dropped, the number of witnesses available for TVSA decreases and BCMA requires more time to reach the τ threshold of beacons. One solution would be to increase T so beacons are sent less frequently. However, we also found that when vehicles travel at high speeds (145km/h), less frequent beacon broadcasts reduce detection capabilities by 30%. To reduce network contention and maintain an acceptable detection level, OBUs could only broadcast witness values for some subset of the beacons heard; for example, only the vehicles that claim a position within their convoy.

B. Impact of Time Synchronization Error on TVSA

When TVSA is deployed in VANETs, a high level of time synchronization may not always be available for OBUs. However, for TVSA to detect malicious activity an OBU needs to determine the difference between its internal time and the internal time of other witnesses. Here we examine how much synchronization is needed and if the clock offset calculation mechanism from Section IV-B can help.

To determine the impact of time synchronization error on TVSA-GS and TVSA-NS, we ran multiple simulations with synchronization errors between $10ns$ and $500ns$. For each simulation of TVSA-GS, we assume the system knows the maximum legitimate synchronization error between OBUs and could therefore adjust the acceptable broadcast time estimate error Δ. As expected, a larger Δ allows a fixed false positive rate of less than 0.5%, but the true positive rate varies. For TVSA-NS, Δ remains at $\pm 10ns$ independent of synchronization error without increasing the false positive rate above 0.5% or decreasing the true positive rate.

Figure 8 presents the detection capabilities of TVSA with and without offset calculation for different amounts of synchronization error between OBUs for attack #3. As anticipated, TVSA with the offset calculation mechanism can detect location claims that break vehicle sequence even when OBUs are not synchronized. TVSA-NS outperforms TVSA-GS for attack #3 on the highway by 75% (90% versus 15%) and 80% in the city (80% vs. 0%) when synchronization is no longer available. We also ran a simulation with a synchronization error of $1ms$ and found TVSA-NS achieves the same detection capabilities.

Here we only present the results for attack #3. For attacks #1 and #2, BCMA can detect the attacker's inconsistent claims (i.e., the attacker's claimed location is forced to move in a sweeping motion or travel in the opposite direction than it claims) before τ beacons are received, so TVSA is most susceptible to attack #3. The addition of offset calculation only improves detection of attack #1 and #2 by 5-10%.

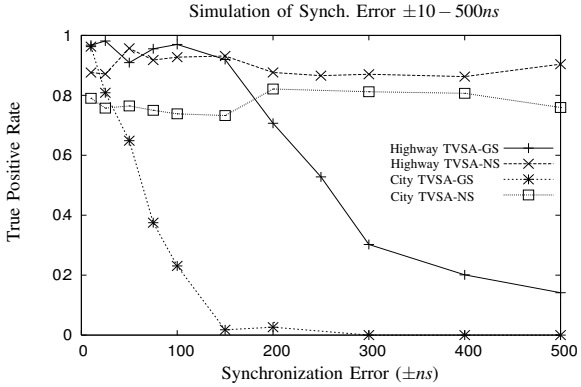

Fig. 8. True Positive Rate vs. Sync. Error for TVSA

The results in this section show that TVSA can operate even if time synchronization between OBUs is not possible. If GPS signals are blocked or RSUs are too costly to install, TVSA-NS can use offset calculation to detect when vehicles claim a location that violates the true vehicle sequence.

C. BCMA + TVSA Overhead Analysis

We analyze the additional overhead of BCMA and TVSA with respect to communication and computation.

BCMA has no additional communication overhead and only a few kilobytes of storage overhead. The CCW application already requires OBUs to broadcast periodic beacons several times a second. All BCMA requires is a counter to track how many beacons the OBU has heard from a specific sender. Even in dense traffic with hundreds of vehicles in range, the storage and management of these counters will require a few kilobytes of memory and minimal processing power. TVSA introduces communication overhead in the form of witness values in the packet and computation overhead when OBUs need to estimate synchronization errors (see Section IV-B). Here we examine the average number of witness values included in beacons for varying highway densities and the amount of computation needed to calculate OBU clock offsets when GPS or RSUs are not available for time references.

Figure 9 indicates the average number of witness values each OBU included in a beacon for varying traffic densities in our simulations. The average number of beacons grows linearly with the increase in traffic density. Each witness value must include an identifier (the original beacon sender's public key or a hash of it), the reception time, and the relative distance. These three items will add 44 bytes or 76 bytes if the hash or entire public key is included. Fourteen witness values add 1KB to a packet. To prevent a high overhead when traffic becomes congested, vehicles could limit their beacons to only include witness values for vehicles that claim to be part of a smaller convoy (e.g., convoy(100m,30) where the OBU considers only

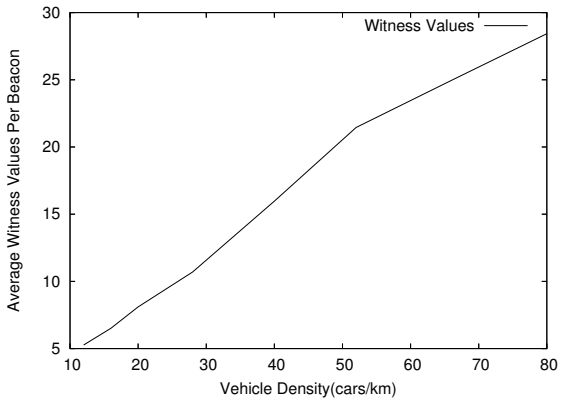

Fig. 9. The Average Number of Witness Values vs. Traffic Density

the 30 closest OBUs within a 100 meter radius, above and below the threshold τ for BCMA). Under such a mechanism, the interval between beacon broadcasts can remain high while the overhead does not exceed a fixed predetermined value of 30 beacons or 2.2KB.

When an OBU does not have a reliable source for time synchronization, the OBU must calculate each witness's clock offset. If an OBU hears beacons from n vehicles, the OBU may check the claims in all n beacons. For each verification, at worst, the OBU must calculate the error of all $n-1$ witnesses using the other $n-2$ witness values. The maximum amount of computation is $n(n-1)(n-2)$ calculations. An $O(n^3)$ algorithm seems computationally expensive but the number of witnesses a vehicle encounters is limited to a small value at any given time. Even with 50 witnesses that all hear each other (thus the full $n(n-1)(n-2)$ operations, TVSA-NS only requires 117,600 efficient calculations. Assuming each operation takes 10 cycles on a 400Mhz machine, TVSA calculation requires 3ms, or roughly half the time to perform 1 ECDSA verification [20]. Within VANETs, the majority of the computation overhead is related to the verification and generation of digital signatures, not our mechanisms.

We have found that our BCMA and TVSA mechanisms can accurately detect when vehicles try to claim a location with opposing traffic or a location that deviates from the real sequence of vehicles. In addition, TVSA-NS provides the same level of detection capabilities independent of time synchronization error. Given the efficiency of BCMA and TVSA, we observe that they are practical even for low-cost OBUs.

VI. Related Work

Several articles were published on general VANET security [10], [18], [20], [21], [27]. These articles frame the general VANET security challenges but consider identification as the most important property for VANET security and do not address the physical properties we present in this paper.

One of the central security challenges in VANETs is establishing trust among vehicles, taking into account their real-world life cycle. Some works suggest using government entities as certificate authorities to help identify valid vehicles [10], [20]. Other work suggests using certificates and TESLA [19], an authentication scheme based on symmetric key operations and delayed key exposure, to establish trust in VANETs [9]. Such an approach reduces security overhead, a TESLA authenticator is 80 bits and takes much less time to verify than a 512 bit ECDSA authenticator. However, TESLA does not provide non-repudiation. In another vein, researchers suggested group signatures or "entanglement" mechanisms to provide privacy, however, they only present a high-level description and did not work out the details [18]. Several recent works have addressed the issue of trust establishment in ad hoc networks, but these mechanisms are not applicable to vehicular networks because they are designed for human interactions [5].

Several researchers considered defenses against specific VANET attacks. For example, Xiao et al. [26] study the detection of Sybil attacks through analysis of radio communication.

Our TVSA protocol makes use of standard distance-bounding protocols, as pioneered by Brands and Chaum [2]. Related to our approach is the area of position verification, where researchers suggested extensions to distance bounding in ad hoc and sensor networks to help nodes determine their position [4], [12]–[15]. In contrast, localization allows a node to obtain an estimate of the position of another node [2], [3], [23]. However, these works do not apply to VANETs since they only determine if the node in question is within a given radius [2], [23] or rely on trusted infrastructure (which may not be available) to make measurements [3].

VII. Conclusion

VANETs are on the verge of wide-spread deployment. Initially, non-safety critical applications will be deployed, such as road toll payments. However, in the very near future the wireless communication capability of DSRC will enable cars to exchange safety-critical information to reduce the frequency and severity of accidents. As we show in this paper, the security mechanisms proposed by the IEEE 1609.2 standard are insufficient to cover the security requirements of many applications; in particular the properties of Convoy Member Authentication (CMA) and Vehicle Sequence Authentication (VSA) that we propose turn out to be crucial for defending against several attacks. We propose new mechanisms for achieving CMA and VSA. Our approaches enable us to secure many position dependent safety applications, which will likely drive the deployment of VANETs. Simulation results show that our mechanisms effectively detect spurious and malicious location claims on highways where vehicle density is decreased. However, further work is required to develop techniques with better detection accuracy within urban environments.

VIII. Acknowledgments

We would like to thank Fan Bai for insightful discussion and feedback on this work. This research was supported in part by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office, and by General Motors

through the GM-CMU Collaborative Research Laboratory. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express or implied, of ARO, CMU, or GM.

REFERENCES

[1] F. Bai, T. Elbatt, G. Hollan, H. Krishnan, and V. Sadekar. Towards characterizing and classifying communication-based automotive applications from a wireless networking perspective. In *Proceedings of IEEE Workshop on Automotive Networking and Applications (AutoNet)*, Dec. 2006.

[2] S. Brands and D. Chaum. Distance-bounding protocols. In T. Helleseth, editor, *Advances in Cryptology – EUROCRYPT '93*, 1993.

[3] S. Capkun, M. Cagalj, and M. Srivastava. Secure localization with hidden and mobile base stations. In *Proc. of INFOCOM*, Mar. 2006.

[4] S. Capkun and J. Hubaux. Secure positioning of wireless devices with application to sensor networks. In *Proc. of INFOCOM*, Mar. 2005.

[5] S. Capkun, J.-P. Hubaux, and L. Buttyan. Mobility helps security in ad hoc networks. In *Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, June 2003.

[6] J. Duffy. U.S. pitches wireless highway safety plan. *Network World*, Nov. 2005.

[7] J. Elson, L. Girod, and D. Estrin. Fine-grained network time synchronization using reference broadcasts. In *Proceedings of Symposium on Operating Systems Design and Implementation (OSDI)*, Dec. 2002.

[8] P. Golle, D. Greene, and J. Staddon. Detecting and correcting malicious data in vanets. In *Proceedings of the Workshop on Vehicular Ad Hoc Networks (VANET)*, 2004.

[9] Y.-C. Hu and K. P. Laberteaux. Strong VANET security on a budget. In *Proceedings of Workshop on Embedded Security in Cars (ESCAR)*, 2006.

[10] J.-P. Hubaux, S. Capkun, and J. Luo. The security and privacy of smart vehicles. *IEEE Security & Privacy magazine*, 2004.

[11] IEEE. 1609.2: Trial-use standard for wireless access in vehicular environments-security services for applications and management messages. IEEE Standards, 2006.

[12] L. Lazos and R. Poovendran. SeRLoc: Secure range-independent localization for wireless sensor networks. In *Proceedings of ACM Wireless Security Workshop (WiSe)*, Oct. 2004.

[13] L. Lazos, R. Poovendran, and S. Capkun. ROPE: Robust position estimation in wireless sensor networks. In *Proceedings of IPSN*, Apr. 2005.

[14] Z. Li, W. Trappe, Y. Zhang, and B. Nath. Robust statistical methods for securing wireless localization in sensor networks. In *Proceedings of IPSN*, Apr. 2005.

[15] D. Liu, P. Ning, and W. K. Du. Attack-resistant location estimation in sensor networks. In *Proceedings of IPSN*, Apr. 2005.

[16] M. A. Lombardi, L. M. Nelson, and A. N. Novick. Time and frequency measurements using the global positioning system. *Cal Lab: The International Journal of Metrology*, July–September 2001.

[17] National Highway Traffic Safety Administration. 2005 state traffic data. http://www-nrd.nhtsa.dot.gov/pdf/nrd-30/NCSA/TSF2005/StateTrafficData05.pdf, Sept. 2006.

[18] B. Parno and A. Perrig. Challenges in securing vehicular networks. In *Proceedings of Workshop on Hot Topics in Networks (HotNets-IV)*, Nov. 2005.

[19] A. Perrig, R. Canetti, D. Tygar, and D. Song. Efficient authentication and signature of multicast streams over lossy channels. In *Proceedings of the IEEE Symposium on Research in Security and Privacy*, May 2000.

[20] M. Raya and J.-P. Hubaux. The security of vehicular ad hoc networks. In *Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, Nov. 2005.

[21] M. Raya and J.-P. Hubaux. Securing vehicular ad hoc networks. *Journal of Computer Security, Special Issue on Security of Ad Hoc and Sensor Networks*, 2007.

[22] A. K. Saha and D. B. Johnson. Modeling mobility for vehicular ad hoc networks. In *Proceedings of the Workshop on Vehicular Ad Hoc Networks (VANET)*, 2004.

[23] N. Sastry, U. Shankar, and D. Wagner. Secure verification of location claims. In *Proceedings of the ACM Workshop on Wireless Security (WiSe)*, Sept. 2003.

[24] VINT Project, University of Berkeley/LBNL. NS-2:network simulator. http://www.isi.edu/nsnam/ns/.

[25] D. Wagner. Resilient aggregation in sensor networks. In *Proceedings of the Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, 2004.

[26] B. Xiao, B. Yu, and C. Gao. Detection and localization of sybil nodes in vanets. In *Proceedings of Workshop on Dependability Issues in Wireless Ad Hoc Networks and Sensor Networks (DIWANS)*, Sept. 2006.

[27] M. E. Zarki, S. Mehrotra, G. Tsudik, and N. Venkatasubramanian. Security issues in a future vehicular network. In *Proceedings of European Wireless*, Feb. 2002.