

# Message Sieving to Mitigate Smart Gridlock Attacks in V2V

Siddharth Dongre

Rochester Institute of Technology  
Rochester, New York, USA  
sd4767@g.rit.edu

Hanif Rahbari

Rochester Institute of Technology  
Rochester, New York, USA  
rahbari@mail.rit.edu

## ABSTRACT

Growing deployment of vehicle-to-vehicle (V2V) communications is expected to significantly increase the volume of Basic Safety Messages (BSM) in highways and dense roads. Computational overhead of verifying the integrity of BSMs will therefore be high while current V2V equipment can process only a limited number of BSMs per second. As a result, critical BSMs carrying vital information may fail to be processed on time, creating unsafe outcomes. In this paper, we expose this vulnerability, discuss critical scenarios, develop novel attacks that exploit this vulnerability, and propose a sieving technique to mitigate these verification gridlock attacks. We show on a USRP testbed that our proposed sieving mechanism to counter sophisticated attackers who exploit this vulnerability achieves 80% accuracy at SNR greater than 6 dB, effectively mitigating the attack.

## CCS CONCEPTS

• **Hardware** → **Digital signal processing**; • **Security and privacy** → *Mobile and wireless security*.

## KEYWORDS

Connected vehicle security, PHY-layer authentication, USRP testbed

### ACM Reference Format:

Siddharth Dongre and Hanif Rahbari. 2021. Message Sieving to Mitigate Smart Gridlock Attacks in V2V. In *14th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '21)*, June 28–July 2, 2021, Abu Dhabi, United Arab Emirates. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3448300.3467834>

## 1 INTRODUCTION

In vehicle-to-vehicle (V2V) communications, vehicles communicate with each other to enhance their safety and proximity awareness [5]. Every vehicle periodically broadcasts a Basic Safety Message (BSM) to announce its current state of motion, including its speed, direction of movement, acceleration and location [1]. As such, some BSMs may act as warning messages intended to prevent impending accidents. Recent trends show that the number of V2V-equipped vehicles is increasing [6]. Therefore, these warning BSMs are becoming more crucial as more vehicles are relying on BSMs for their actions. On dense roads, however, transmission of a large number of




Figure 1: Example of a critical scenario in a dense road.

BSMs can create a processing bottleneck at receiving vehicles. This may create potentially harmful outcomes if critical BSMs fail to be processed on time. For example, during peak hours on the I-490 highway outside Rochester, NY, where the likelihood of vehicle collisions is high, the projected number of transmitted BSMs based on I-490's current traffic data [9] can easily exceed the processing capability of the state-of-the-art V2V chipsets, as exemplified in Figure 1.

The processing bottleneck at receiving vehicles is in part due to the verification of digital signatures. The National Highway Traffic Safety Administration (NHTSA) requires using digital signatures with certificates to authenticate the sender of BSMs and verify the integrity of their content [1]. This ensures that BSMs come from a legitimate vehicle, not from a malicious transmitter, and are not tampered with in transit. Although this verification is vital to the security of a V2V system, it unfortunately adds an extra computational overhead on receiving vehicles. This overhead in turn may impact the availability of BSMs at a receiving vehicle. Ideally, each BSM should be processed as soon as it is received. The chipsets used in modern V2V equipment are capable of verifying up to 2500 BSMs per second [6]. We show in this paper that, unfortunately, this upper limit is not sufficient to process all BSMs in a realistically dense road, creating a BSM gridlock state. As a result, the receiver fails to process all of the received BSMs, meaning that it discards an arbitrary subset of BSMs in a given interval, which could lead to an unsafe outcome if the discarded BSMs contain critical messages.

One may consider mitigating this BSM bottleneck vulnerability by applying a BSM filter that discards non-critical BSMs before performing the signature verification. Although this is a promising approach, we show that a naive filter will perform poorly once an attacker tries to proactively bypass it, leaving the receiver unable to process all critical BSMs. We further identify attack scenarios in different settings with severe consequences. Therefore, the filter requires sophistication to be effective against not just simple attacks, such as simply sending a burst of malicious BSMs, but also smart attacks that try to circumvent a naive filter. Designing an effective filter requires addressing two main challenges. First challenge is that it needs to differentiate critical BSMs from non-critical ones quickly and reliably. Second, the filter needs to be robust enough to distinguish malicious BSMs from the ones sent by benign vehicles.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

WiSec '21, June 28–July 1, 2021, Abu Dhabi, UAE

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8349-3/21/06...\$15.00

<https://doi.org/10.1145/3448300.3467834>

**Table 1: Urban setting - Possible consequences of BSM verification bottleneck.**

Type of BSM (%)	One vehicle discards:		Multiple vehicles discard:	
	1 BSM	Multiple BSMs	1 BSM	Multiple BSMs
<b>FCW and EEBL (44%)</b>	A rear-end collision	A rear-end collision	<b>*Multiple rear-end collisions</b>	Multiple rear-end collisions
<b>IMA (23%)</b>	A T-bone collision	A T-bone collision	<b>*Multiple T-bone collisions</b>	Multiple T-bone collisions
<b>LTA (9%)</b>	A T-bone collision	A T-bone collision	Multiple collisions	Multiple collisions
<b>DNPW (10%)</b>	A head-on collision	A head-on collision	Multiple head-on collisions	Multiple head-on collisions
<b>BS / LCW (14%)</b>	A sideswipe	A sideswipe	<b>*Multiple sideswipes</b>	Multiple sideswipes

To address these challenges, we first need to quantify the most critical BSMs and understand the significance of these BSMs based on the severity of the scenarios each of these BSMs is intended to prevent. We also identify specific locations as representative examples for such critical scenarios. This can further be used to evaluate the effectiveness of the filter. For this purpose, we use the U.S. Department of Transportation’s (DoT) vehicle usage data that is then fed into the advanced traffic simulation software PTV Vissim [3]. Simulated traffic data is used to study the likelihood of unsafe outcomes in those scenarios.

To identify non-critical BSMs coming from irrelevant vehicles (e.g., vehicles in the opposite direction in a split-lane highway), we employ a physical (PHY) layer Doppler shift estimation technique to accurately determine the relative velocity of a transmitting vehicle with respect to the receiver to be able to sieve the non-critical BSMs. Most V2V implementations utilize Dedicated Short-Range Communications (DSRC) protocol [11]. We apply our estimation technique on IEEE 802.11p frame preamble taking into account hardware imperfections that may lead to an erroneous estimation. To evaluate and enhance the performance of our filter, we implement a simple gridlock attack that only blindly transmits bursts of fake BSMs to exploit the bottleneck vulnerability, and a smart attack that also attempts to circumvent a basic version of our filter. Consequently, we strengthen our filter to be effective against such smart attacks.

To summarize, our contributions are as follows:

(1) Using realistic traffic scenarios, with and without the presence of an attacker, we bring to light that in dense roads at least 10% BSMs are at the risk of not being processed by the receiver.

(2) We develop a simple attack and show that it can increase the probability of discarding all types of BSMs by at least 18%. Additionally, we develop a smart gridlock attack and show that it can reduce the accuracy of a naive BSM sieving mechanism to less than 50%.

(3) We develop the first BSM filtering/sieving mechanism against not only simple but also smart gridlock attacks. Our USRP-based experiments show that our improved PHY-layer filter can achieve 80% accuracy, allowing a receiver to handle significantly higher number of BSMs in gridlock than what current V2V chipsets can.

In the remainder of the paper, we first provide an overview of related work in Section 2. In Section 3 we describe the BSM verification bottleneck vulnerability, and then describe our threat model in Section 4. Critical attack scenarios are introduced in Section 5 and we present our proposed filtering technique in Section 6. Finally, we experimentally evaluate the performance of our filter in Section 7 before concluding the paper in Section 8.

## 2 RELATED WORK

In this section we review existing work related to the performance of DSRC in dense roads, and the techniques to identify malicious BSMs. In [4], the authors develop an analytical model to study the reliability of BSM delivery under dense and highly dynamic environments and show that BSM packet delivery delay increases with vehicle density. The impact of urban environments’ density on the throughput of DSRC traffic has been studied in [2] showing that dense environments severely impact the throughput because of large number of packet collisions. However, these works do not consider the bottleneck due to V2V chipsets’ processing limitations.

To detect and filter BSMs sent by malicious vehicles, the authors in [12] propose a reputation management scheme. However, checking the reputation of a vehicle still requires authenticating it via digital signatures, which takes processing time. Sun *et al.* in [13] present a scheme which estimates the angle-of-arrival and frequency-difference-of-arrival of the received BSMs to cross check the location and motion claim of a vehicle. Their method relies on the availability of valid reflecting surfaces in the environment creating multi-path signals. Such reflecting surfaces may not always be available. Moreover, it needs to verify the BSM digital signature to validate the velocity claim of the transmitting vehicle. Our proposed method does not require verification of the BSM digital signature to validate the velocity information of the transmitter.

## 3 BSM VERIFICATION BOTTLENECK

*Critical BSMs.* We start by reviewing the BSMs that are considered to be of highest importance by the NHTSA. These BSMs can potentially help to prevent the top-10 pre-crash scenarios in 49% of all crashes in the U.S. [1]. These BSMs and the type of collisions they intend to prevent are as follows:

- Forward Collision Warning (FCW) and Emergency Electronic Brake Lights (EEBL): Rear-end collisions
- Intersection Movement Assist (IMA) and Left Turn Assist (LTA): T-bone collisions
- Do Not Pass Warning (DNPW): Head-on collisions
- Blind Spot / Lane Change Warning (BS / LCW): Sideswipe collisions when changing lanes

Failure to process these BSMs in a dense road would negate the primary objective of V2V, potentially creating unsafe consequences. We quantify the importance of each BSM type using the percentage  $\%_0$  to denote the ratio of the number of specific collisions intended to be prevented by a given type of BSM (as listed above) to the number of all collisions. We calculated  $\%_0$  based on the estimates made by NHTSA in 2017 [1], assuming that each collision occurs

Table 2: Highway setting - Possible consequences of BSM verification bottleneck.

Type of BSM (%)	One vehicle discards:		Multiple vehicles discard:	
	1 BSM	Multiple BSMs	1 BSM	Multiple BSMs
FCW and EEBL 164%	A rear-end collision	A rear-end collision	*Multiple rear-end collisions	Multiple rear-end collisions
DNPW 116%	A head-on collision	A head-on collision	Multiple head-on collisions	Multiple head-on collisions
BS / LCW 120%	A sideswipe	Single sideswipe	*Multiple sideswipes	Multiple sideswipes

when a vehicle fails to process the corresponding warning BSM. We then consider four scenarios where one or multiple BSMs are discarded on (a) one, or (b) multiple vehicles. Tables 1 and 2 show the importance of each BSM type (based on the outcomes they are expected to prevent in each scenario for urban and highway settings, respectively, highlighting critical scenarios with severely unsafe outcomes (marked with \*). In Section 5 we will present attacks that can be performed on these critical scenarios.

To study the impact of a dense, straight road on the BSM processing bottleneck, we define  $\lambda$  as the average number of BSMs received per second by any vehicle on that road, where  $v$  is average speed,  $s$  is average car spacing in seconds,  $n$  is number of lanes per direction,  $\mu$  is average BSM generation rate by a single vehicle (usually 10 per second), and  $r$  is the communication range of vehicles. Note that the area of the rectangular communication range of a given vehicle on a straight road is  $2r \cdot s$ . One can verify by geometric and dimensional analysis (also shown analytically in [9]) that  $\lambda$  increases linearly with  $v$ ,  $\mu$ ,  $n$ , and  $r$ , and decreases with  $s$ , as follows

$$\lambda = \frac{2 \cdot \mu \cdot n \cdot r \cdot v}{s} \quad (1)$$

As an example, consider the I-490 highway mentioned earlier and a DSRC-based V2V, where the average vehicle speed is 60 mph (27.5 m/s), vehicle spacing is a typical 5B [9], and  $\mu = 100$  on a three-lane highway with two-way traffic. The number of BSMs generated in this environment would be approximately 3600 per second resulting in about 10% around 30% of these BSMs getting discarded at the chipsets with maximum BSM processing rate (e.g., Qualcomm 9150) [14]. In a dense highway, the likelihood of a BSM to be a warning message is usually high. As a result, there is a high likelihood of such a scenario to result in a harmful outcome.

#### 4 THREAT MODEL

We consider two attack types that try to exploit the bottleneck vulnerability, a simple and a smart attack. They differ in their method of attack, but both share a common goal creating disruption in the natural flow of traffic and affecting as many vehicles as possible, potentially creating a collision. The attacker generates and transmits fake BSMs while it does not need to generate valid signatures as it only aims to throttle the receiver's processor by forcing excessive signature verifications and does not care if verifications fail. We argue that since the attacker is trying to create a disruption in traffic flow, being a part of that traffic flow is counter intuitive. Hence, we assume the attacker is stationary as opposed to a mobile attacker who is moving along target vehicles. They can arbitrarily choose

their location of attack using publicly available traffic information, such as annual average daily traffic (AADT) data [9] to select the optimum location that can create most disruption. We further assume that the attacker has an idea of the average velocity by utilizing BSMs sent by target vehicles indicating their instantaneous velocity or inferring the speed limit of the road they are targeting.

In the simple attack, the attacker tries to further throttle the bottleneck at vehicles by transmitting a burst of BSMs in every BSM interval (i.e., 100ms). It uses the default 802.11p parameters to transmit these BSMs without any modification to the PHY-layer frames. The burst size is less than 10% of the total number of BSMs in the environment (which can be estimated as shown in Section 3), allowing it to remain undetected. We confirmed an attacker can send up to 40 BSMs in an interval through experiments performed on commercially used V2V chipset (Qualcomm 9150) on Cohda Wireless MK6C kits [14].

In the smart attack, the attacker further tries to circumvent a naive mitigation technique by adjusting the PHY-layer parameters (e.g., carrier frequency offset) of the BSMs it generates. Its intention is not to be detected as a non-participant member of the traffic flow. We develop the smart attack in further details in Section 6.3 after describing our basic filtering mechanism.

#### 5 GRIDLOCK ATTACK SCENARIOS

In this section we describe specific critical scenarios from the ones identified in Section 3 where an attacker can launch a destructive BSM gridlock attack. Denote the probability of a specific BSM warning getting discarded at a receiving vehicle as

$$P_{\text{disc}} = \frac{\#_{\text{disc}}}{\#} \quad (2)$$

where  $\#$  is the total number of all BSMs received by that vehicle at a given interval and  $\#_{\text{disc}}$  is the number of vehicles sending a specific critical BSM in the same interval. We derive this relation by recognizing that the events of discarding specific BSM warnings sent by different vehicles are identically and independently distributed. The following scenarios are specific instances from selected locations. Similar scenarios may exist in other locations, but an exhaustive study of such locations is out of the scope of this paper. The main goal behind analyzing these scenarios is to show how  $P_{\text{disc}}$  changes in different scenarios, and that a receiver discards at least 0% of the BSMs ( $P_{\text{disc}} > 0$ ) even when there is no attacker. We performed simulations in MATLAB using traffic data obtained from PTV Vissim traffic simulator [3].

##### 5.1 Scenario 1 - FCW

Consider a stretch of MD-295, known as Baltimore-Washington Pkwy. It is notorious for high traffic density and low vehicle spacing

of 175s. We calculated the vehicle spacing based on the AADT data obtained from Maryland's DoT [8]. A snapshot of the simulated traffic for this scenario is shown in Figure 2, where each small colored box on the road represents a vehicle. Applying (2) to the vehicle density data gathered from each simulation run, we observe that in a traffic column with 20-25 vehicles, an FCW is dropped at a receiving vehicle in the first interval with an average probability of  $\frac{1}{186}$ . Here,  $\frac{1}{186}$  changes with vehicle density in each run but  $\leq 2$  remains constant as we assume 2-lane traffic here, meaning that only the two leading vehicles will be able to immediately send an FCW once they detect a collision in front of them. By applying (2) recursively, where  $\tau$  increases in each interval as more vehicles have received the FCW, it turns out that it takes an average of 4.5 intervals for an FCW to reach all the vehicles in the column. With current hardware, an attacker can transmit bursts of 40 BSMs per interval to exacerbate the gridlock while remaining undetected. This will increase the likelihood of dropping BSMs by almost 40% and increase the BSM processing delay by 2 more intervals. This delay can lead to collisions in cases where the vehicles are tightly packed as is the case with this scenario.

### 5.2 Scenario 2 - LCW

We perform the same analysis on a stretch of road on the I-490 highway outside Rochester, NY. Based on AADT data gathered from New York's DoT we determined the average vehicle spacing to be 175s [9]. This scenario focuses on vehicles that are sending LCW when changing lanes. As the LCW will only be sent out by the vehicle attempting the lane change, we set  $\tau = 1$  in the first interval. It turns out that the vehicles discard an LCW with an average probability of  $\frac{1}{343}$  and that increases by 18% in the presence of an attacker sending bursts of 40 BSMs per interval.

### 5.3 Scenario 3 - IMA

Now we look at an urban intersection scenario involving vehicles that rely on the IMA message to avoid collisions. We considered the intersection between MD-97 (Georgia Avenue) and MD-193 [10]. We look at the case with a single vehicle attempting to cross the intersection. There is one vehicle sending IMA warnings, i.e.,  $\tau = 1$ . Traffic simulation showed over 300 vehicles in a 1 km radius from the intersection. We determined that vehicles approaching the intersection would discard the IMA with an average probability of  $\frac{1}{1910}$  that increases by 37% when an attacker sends a burst of just 30 BSMs in an interval.

## 6 PROPOSED BSM FILTERING TECHNIQUE

We propose a filtering mechanism at the PHY layer, before any signature verification at upper layers, based on the relative velocity of the sender of a BSM to determine the importance of that BSM. The benefit of this approach is that it incurs negligible overhead while it reduces the signature verification overhead, allowing the receiver to verify only the BSMs that pass through the filter.

### 6.1 Doppler Shift Estimation

The relative velocity can be estimated by measuring the carrier frequency offset (CFO) caused by the Doppler shift at the PHY-layer. However, the CFO is a result of not only Doppler shift  $\Delta f_D$

due to relative motions of the transmitting and receiving vehicles, but also the difference in the operating frequencies of two radio oscillators due to oscillator imperfections. To accurately estimate the relative velocity, we need to separate  $\Delta f_D$  from the total CFO. To start, we use a reliable method to estimate the overall CFO  $\Delta f$  using the frame preamble, as explained in [11]. The preamble in 802.11p is a periodic signal. Two cycles, each 8 samples long with sampling period  $T_s$ , can be used to create a sequence. Let  $A_n$  be the  $n$ th sample of the sequence where  $n = 1, \dots, 2L$ . As the sequence is made up of identical cycles  $A_{n+2L} = A_n$ . Note that applying a resultant CFO of  $\Delta f$ , the phase of  $A_n$  is offset by  $\phi_n = 2\pi \Delta f n T_s$ . Therefore, to estimate phase offset  $\phi$ , the receiver first multiplies the conjugate of  $A_{2L}$  with  $A_8$  to obtain  $B_8 = 1 \cdot e^{-j\phi}$ ,

$$B_8 \stackrel{\text{def}}{=} A_8 A_{2L}^* = j A_8^* A_{2L}^{2\pi \Delta f 8 T_s} = j A_8^* A_{2L}^{2\pi \Delta f 8 T_s} \quad (3)$$

Then it accounts for noise (not shown) by adding  $B_8$  to obtain

$$\hat{z}_8 = \sum_{n=0}^{L-1} B_{8+2n} \quad (4)$$

where  $\hat{z}_8$  denotes the phase of a complex number. The estimated CFO is thus given by,

$$\hat{f}_5 = \frac{\hat{z}_8}{2\pi 8 T_s} \quad (5)$$

The overall CFO is also the sum  $\Delta f = \Delta f_D + \Delta f_{osc}$ , where  $\Delta f_{osc}$  is the sampling clock drift related to oscillator imperfections. However, we are only interested in  $\Delta f_D$ . Using the method proposed in [12], we can estimate sampling clock drift in mobility scenarios. Thus, the estimated Doppler shift due to mobility is given by  $\Delta f_D = \hat{f}_5 - \Delta f_{osc}$ . Finally, we estimate the relative velocity between transmitter and receiver as,

$$E_{A_4} = \frac{v}{c} \frac{f_c}{2} \quad (6)$$

where  $v$  is the speed of light ( $c = 3 \times 10^8$  m/s),  $f_c$  is the carrier frequency (5.9 GHz), and  $E_{A_4}$  is the relative velocity. It is negative when transmitter is going away from the receiver, and positive when the transmitter is moving toward the receiver.

Figure 2: Snapshot of PTV Vissim showing simulated vehicle columns on Baltimore-Washington Pkwy.

## 6.2 Basic Filter

After determining the relative velocity, the receiving vehicle filters the BSM based on a set of filtering rules that classify the transmitting vehicle into three types: (1) vehicle in the opposite direction, (2) vehicle in the same direction moving away from the receiver, and (3) vehicle in the same direction moving toward the receiver. The filter is active only when the receiver is moving, and the road is dense. The receiver observes the average number of BSMs it receives per second to decide whether to activate the filter. In a gridlock situation, the receiver is only concerned with BSMs sent by vehicles moving in the same direction and are getting closer to the receiver, such as FCW sent by stopping vehicles ahead (see Section 5.1). Therefore, BSMs from Type 1 and Type 2 vehicles can be safely filtered.

To check whether the transmitter is in the opposite direction, we compare the magnitude of the relative velocity against a threshold. A transmitter moving in the same direction will have a maximum magnitude of relative velocity equal to  $v_{rel} = v_{D} + v_{l}$ , where  $v_{D}$  is the receiver's speed and  $v_{l}$  is the speed limit. Thus, a relative velocity of higher magnitude indicates the vehicle is moving in the opposite direction (Type 1). We also know that the sign of the relative velocity  $v_{rel}$  indicates whether the transmitter is moving toward or away from the receiver. Let us specify how aggressive the filtering is. Accordingly, an active filter classifies transmitting vehicles as follows:

- Type 1: If  $|v_{rel}| > v_{D} + v_{l}$ , then the transmitter is moving in the opposite direction.
- Type 2: If  $|v_{rel}| < v_{D}$  and  $v_{rel} < 0$ , then the transmitter is going away from the receiver.
- Type 3: If  $|v_{rel}| < v_{D}$  and  $v_{rel} > 0$ , then the transmitter is moving closer to the receiver.

## 6.3 Smart Attack against Basic Filter

Because the simple attacker is assumed to be stationary ( $v_{rel} = D$ ), the BSMs sent by such an attacker will get filtered easily at the receiver. However, if the attacker is smart and is able to apply an artificial CFO such that it is able to emulate a moving vehicle, its transmitted BSMs can bypass the basic (naïve) filter above.

Consider the receiving vehicles to be moving at an average velocity  $E$ . We consider the attacker to be parked near a location it has determined to be the most optimum to create disruption (e.g., see the locations in Section 5). The attacker emulates ghost vehicle moving at a velocity  $E$ . To do this, it transmits a BSM after applying an artificial CFO  $\Delta f$  equal to the Doppler shift corresponding to velocity  $E$ . The receiver would then estimate  $f = f_0 + \Delta f$ , where  $\Delta f$  is the CFO due to the motion of receiver. Thus,  $f_{rel} = E \cdot f_0$ . The attacker selects the value of  $\Delta f$  from a range such that  $f_{rel}$  satisfies the definition of a Type 3 vehicle at the receiver and bypasses the filter. To be able to detect multiple vehicles on the road, the attacker can set the magnitude of the speed limit, as a majority of the receiving vehicles will be moving at this velocity. With the basic filter, the receiving vehicles cannot differentiate between real motion and an artificially applied CFO.

Applying Artificial CFO. Consider the sequence containing  $2L$  samples of two cycles of a BSM frame preamble at the transmitter

generated at sampling period  $T_s$ . Doppler shift due to velocity  $E$  is given by,

$$\Delta f = \frac{E}{c} f_0 \quad (7)$$

The phase offset at each sample of the sequence is given by,

$$\phi_l = 2\pi \Delta f l T_s \quad (8)$$

where the vector  $\mathbf{1} = [1 \dots 1]^T$ . Then the resulting sequence is,

$$\mathbf{r} = (\mathbf{1} \odot \mathbf{e}^{j\phi_l}) \quad (9)$$

where  $\odot$  represents element-wise vector multiplication. The attacker applies the artificial CFO to the cycles (or the entire preamble) intended to be used for CFO estimation to generate its preamble.

It can be said that the effects of an applied CFO on a transmitted signal are not exactly the same as Doppler effect due to mobility. We analyzed the effects of both and determined that the true CFO due to Doppler effect and an artificial CFO is within 1 Hz of each other. This difference is negligible as CFO tends to be in the order of magnitude of 100 Hz.

## 6.4 Improved Filter

To make the filter sensitive against smart attacks, we take advantage of certain limitations of the attacker. As explained in our threat model, such an attacker needs to be stationary and can rely only on the speed limit of the target road. The target vehicle can perturb its speed in a small amount of time in an attempt to create a change in relative velocity  $v_{rel}$  of transmitter with respect to time  $C$ . The attacker cannot reliably obtain  $v_{rel}$  and  $C$  in time to change its artificial CFO in an attempt to bypass the filter. For a given estimation of  $v_{rel}$  taken from BSMs sent over a time window  $C$  if  $\frac{v_{rel}}{C} > \frac{v_{l}}{C}$ , the BSM was likely to be sent by the attacker.

The challenge is to choose the value of  $\Delta f$  such that it minimizes false positives and false negatives, i.e., it does not filter packets coming from legitimate vehicles due to small changes in velocity. Our PTVissim simulations show that traffic columns tend to change their velocity in unison with each other. We see vehicles increase their velocity from 45 mph to 55 mph within a few seconds and the change in relative velocities among the different vehicles is small. This allows us to set a small enough  $\Delta f$  to minimize false positives and false negatives.

## 7 PERFORMANCE EVALUATION

We perform experiments using two USRP N210s with UBX40 daughterboards that are stationary in a lab environment. We implemented 802.11p preamble for DSRC and the CFO estimation technique in LabVIEW. The Tx (attacker) and Rx (target vehicle) are 4m apart from each other with the Tx operating at 2 dBm transmit power. We model an AWGN channel with varying synthesized noise power to achieve SNR from 5 to 20 dB. To emulate mobility, we leverage applying an artificial CFO on the transmitted frames that mimic the effects of relative velocity. The transmitter sends frames at 10 Hz rate over a period of 500s. These contain different types of benign and malicious packets, each with a different artificial CFO.

For the simple attack, the CFO due to Doppler shift is large in magnitude and negative, since the attacker is not actively trying to bypass the filter. For the smart attack, that CFO is small in magnitude and positive, as the attacker is intelligently trying to bypass

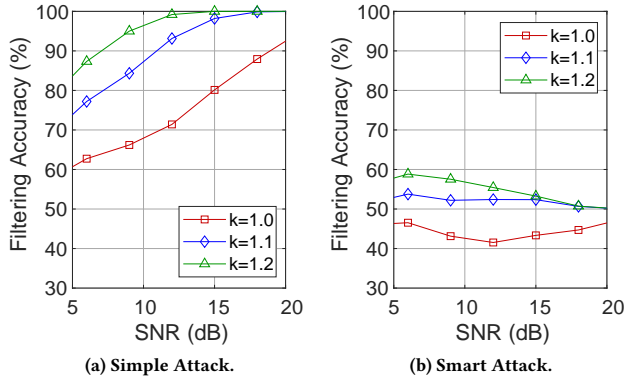


Figure 3: Basic filter performance.

the filter by emulating a vehicle moving at around the same velocity as the target in the same direction. We apply a large CFO for vehicles moving in the opposite direction, while vehicles moving in the same direction have a very small CFO in magnitude. We leverage simulated traffic data obtained from PTV Vissim to generate instantaneous velocity information at each 100 ms time interval and use it to apply specific CFO to specific benign frames.

Figure 3(a) shows the performance of the basic filter defined in Section 6.2 without the improvements. We can see that under the simple attack, the filter is able to achieve accuracy greater than 80% for  $\epsilon = 1^{-1}$  and SNR greater than 10 dB. However, under the smart attack (Figure 3(b)), the basic filter is completely ineffective with accuracy close to 50% overall. This shows that the smart attack is capable enough to effectively bypass the basic filter.

We next study the performance improvements after adapting the filter to withstand smart attacks. The accuracy seems to be independent of  $\epsilon$ , as seen in Figure 4(a), hence we can set it to the lowest possible value of 0.5s for fastest performance. We can see in Figure 4(b) that increasing the value of  $\epsilon$  A4B  $> 3$  when  $\epsilon = 0.5$  has a negative impact on the accuracy because the rate of change in relative velocity is small among benign vehicles. Overall, our filter achieves greater than 80% accuracy for SNR  $\geq 6$  dB, meaning, a receiver can now handle 450 more BSMs per interval by filtering non-critical BSMs. With the improved filter, we substantially neutralize the bottleneck in the scenarios described in Section 5, reducing the probability of BSMs getting discarded to almost zero.

## 8 CONCLUSION

In this paper, we have analyzed traffic scenarios in an urban and highway settings that show the severity of the BSM digital signature verification bottleneck in V2V. We have proposed a BSM filtering mechanism that discards non-critical BSMs on the basis of the relative velocity of the transmitter, estimated using the Doppler shift between two vehicles. We have shown that the improved filtering mechanism is robust against both simple and smart attacks. With the help of realistic traffic simulations and USRP experiments, we have shown that our filter eases the bottleneck with 80% accuracy at SNR  $\geq 6$  dB allowing 450 more BSMs to be handled per interval. In future, we intend to further improve the filter by utilizing different statistical performance metrics.




Figure 4: Improved filter performance under Smart Attack.

## ACKNOWLEDGMENTS

This research was supported by the National Security Agency under Grant Number H98230-19-1-0318. The United States Government is authorized to reproduce and distribute reprints notwithstanding any copyright notation herein. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Security Agency.

## REFERENCES

- [1] National Highway Traffic Safety Administration. 2017. *Federal Motor Vehicle Safety Standards: V2V Communications*. Retrieved March 2, 2021 from <https://tinyurl.com/3nht9awk>
- [2] Stephan Eichler. 2007. Performance Evaluation of the IEEE 802.11p WAVE Communication Standard. In *IEEE Vehicular Technology Conference (VTC)*. Baltimore, MD, USA, 2199–2203.
- [3] PTV Group. 2020. *PTV Vissim*. Retrieved March 2, 2021 from <https://www.ptvgroup.com/en/solutions/products/ptv-vissim/>
- [4] Khalid Abdel Hafeez, Lian Zhao, Bobby Ma, and Jon W. Mark. 2013. Performance Analysis and Enhancement of the DSRC for VANET’s Safety Applications. *IEEE Transactions on Vehicular Technology* 62, 7 (Sept. 2013), 3069–3083.
- [5] Kyusuk Han, Swapna Divya Potluri, and Kang G. Shin. 2013. On Authentication in a Connected Vehicle: Secure Integration of Mobile Devices with Vehicular Networks. In *Proc. ACM/IEEE Int. Conf. on Cyber-Physical Systems (ICCP)*. Philadelphia, Pennsylvania, 160–169.
- [6] Qualcomm Technologies Inc. 2020. *Qualcomm Introduces Comprehensive Platform for RSU and OBU to Further Accelerate C-V2X Global Momentum*. Retrieved March 2, 2021 from <https://tinyurl.com/2rea555>
- [7] Hyunbeom Lee and Jungwoo Lee. 2011. Joint Clock and Frequency Synchronization for OFDM-Based Cellular Systems. *IEEE Signal Processing Letters* 18, 12 (Oct. 2011), 757–760.
- [8] Maryland Department of Transportation. 2019. *Traffic Volume Maps*. Retrieved March 2, 2021 from <https://tinyurl.com/pmyeysne>
- [9] New York Department of Transportation. 2020. *Traffic Data Viewer*. Retrieved May 26, 2021 from <https://www.dot.ny.gov/tdv>
- [10] Hanif Rahbari, Marwan Krunz, and Loukas Lazos. 2016. Swift Jamming Attack on Frequency Offset Estimation: The Achilles’ Heel of OFDM Systems. *IEEE Transactions on Mobile Computing* 15, 5 (July 2016), 1264–1278.
- [11] Murray Slovick. 2017. *DSRC vs. C-V2X: Looking to impress the regulators*. Retrieved March 2, 2021 from <https://tinyurl.com/9pkr9r6m>
- [12] Shen Su, Zhihong Zian, Siyu Liang, Shuang Li, Shasha Du, and Nadra Guizani. 2020. A Reputation Management Scheme for Efficient Malicious Vehicle Identification over 5G Networks. *IEEE Wireless Communications* 27, 3 (June 2020), 46–52.
- [13] Mingshun Sun, Yanmao Man, Ming Li, and Ryan Gerdes. 2020. SVM: Secure Vehicle Motion Verification with a Single Wireless Receiver. In *Proc. ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec’20)*. Linz, Austria, 65–76.
- [14] Cohda Wireless. 2020. *MK6C EVK*. Retrieved March 2, 2021 from <https://cohdawireless.com/solutions/hardware/mk6c-evk/>