Vehicle-to-Nothing? Securing C-V2X Against Protocol-Aware DoS Attacks

Geoff Twardokus and Hanif Rahbari
Rochester Institute of Technology, Rochester, NY, USA
{geoff.twardokus,rahbari}@mail.rit.edu

***Abstract*—Vehicle-to-vehicle (V2V) communication allows vehicles to directly exchange messages, increasing their situational awareness and offering the potential to prevent hundreds of thousands vehicular crashes annually. Cellular Vehicle-to-Everything (C-V2X), with its LTE-V2X and New Radio (NR)-V2X variants in 4G/LTE- and 5G-based C-V2X, is emerging as the main V2V technology. However, despite security protocols and standards for C-V2X, we expose in this paper that its physical (PHY) and MAC layers are not resilient against intelligent, protocol-aware attacks due to the very predictable PHY-layer structure and vulnerable scheduling algorithm used in both LTE-V2X and NR-V2X. We devise two stealthy denial-of-service (DoS) exploits that dramatically degrade C-V2X availability, thereby increasing the chances of fatal vehicle collisions. We experimentally evaluate our attacks on an integrated, hybrid testbed with USRPs and state-of-the-art LTE-V2X devices as well as through extensive simulations, showing that within seconds, our attacks can reduce a target's packet delivery ratio by 90% or degrade C-V2X channel throughput by 50%. We propose, analyze, and evaluate detection approaches as well as mitigation techniques to address the vulnerabilities we expose in the C-V2X PHY/MAC layers, providing direction towards better-secured, resilient 5G C-V2X.**

***Index Terms*—V2V security, denial-of-service, jamming, resource scheduling**

I. Introduction

Cellular Vehicle-to-Everything (C-V2X), the leading technology family for vehicle-to-vehicle (V2V) communications, comprises two cellular technologies: LTE-V2X (3GPP Rel-14/15 [1], [2]) and New Radio (NR)-V2X (3GPP Rel-16/17 [3], [4]). In C-V2X, vehicles use *sidelink* communications in Mode 4 (LTE) or Mode 2 (NR) to directly exchange periodic basic safety messages (BSMs), allowing vehicles to maintain awareness of each others' movements. Among other things, this facilitates collision avoidance in non-line-of-sight (NLOS) scenarios, where a driver or onboard sensors (e.g., LiDAR) cannot perceive an imminent collision. The potential benefits to society are vast: V2V is projected to prevent up to $595,000$ vehicle crashes annually in the U.S. alone, saving the economy as much as $\$71$ billion every year [5]. As LTE-V2X is currently the cornerstone of global C-V2X deployments in smart transportation systems [6]–[8], "5G C-V2X"—*hybrid* deployments wherein LTE-V2X is complemented, not replaced, by NR-V2X—is poised to dominate the V2V space for the foreseeable future [9], [10].

As a safety-critical technology, V2V must be properly secured against malicious attacks. A vehicle moving at high speed may have just a few milliseconds of reaction time in order to avoid a collision, so ensuring the availability of C-V2X communication is especially imperative. Unfortunately, existing V2V security schemes and standards often have one of two key limitations. First, much recent work on V2V security (e.g., [11], [12]) assesses an older, 802.11-based protocol—Dedicated Short Range Communication (DSRC) [13]—and those works which do examine C-V2X (e.g., [14], [15]) often focus their attention on new and unique features of NR-V2X or its upper-layer security. However, it is critical that NR-V2X be considered alongside LTE-V2X because, in addition to their expected long-term coexistence, NR-V2X has inherited many physical (PHY) and MAC-layer elements from LTE-V2X. Second, existing works which have examined security in LTE-V2X (e.g., [16], [17]) generally address confidentiality or integrity concerns, leaving *availability* far less studied. As more and more vehicles rely on V2V, further study of C-V2X availability–in both LTE- and NR-V2X–is essential to ensure the safety and security of future roadways.

LTE and NR sidelink frames comprise rigid time-frequency grids with fixed locations for control and data channels. This structure enables certain beneficial features like simultaneous transmissions and increased spectral efficiency, but it also makes it easy to accurately predict where in the grid certain transmissions (e.g., from a particular vehicle) will occur. For BSMs, it induces far more precise periodicity than in DSRC, so one can precisely predict the times at which a vehicle will transmit its BSMs. Both LTE- and NR-V2X use an autonomous scheduling protocol called *semi-persistent scheduling* (SPS) [18], [19]. SPS is autonomous because sidelink is expected not to rely on eNBs or gNBs for resource allocation. As a sensing-based algorithm, SPS assumes the periodic nature of BSMs and uses past observations to predict future channel usage, allowing a vehicle to identify time-frequency resources for its transmissions that are not likely to be occupied by other vehicles. However, in dynamic scenarios, packet collisions still occur (as shown in [20], [21]) because SPS cannot entirely prevent vehicles from selecting the same resources. This can be exacerbated by an intelligent attacker. Unfortunately, such a denial-of-service (DoS) attack is difficult to detect because its effects may be masked by those benign collisions.

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.

In this paper, we first illuminate the above vulnerabilities by devising and experimentally demonstrating two stealthy, "protocol-aware" DoS attacks and then take steps to detect and counter them. Our attacks have catastrophic effects, allowing attackers their choice of silencing particular vehicles at will or crippling C-V2X system performance as a whole by depriving all vehicles of up to 50% of available spectrum resources.

In *targeted sidelink jamming*, an attacker takes advantage of the rigid PHY-layer structure of LTE-V2X and the precise periodicity of BSMs to predict the arrival of BSMs from a specific target vehicle and jam them. The attacker only needs to analyze one out of every $5-15$ messages sent by its target and jams only a small portion (up to 10% in practical systems) of each of the target's subsequent BSMs. Using USRPs to attack state-of-the-art LTE-V2X devices based on Qualcomm chipsets [22], we achieve a very high degree of accuracy and up to a 93% reduction in packet delivery ratio for the target vehicle. We also show that the effects of this attack cannot be easily distinguished from packet losses due to SPS conflicts, especially under heavier channel loads. To address this, we propose and validate a superior detection technique based on unsupervised cluster analysis.

In *sidelink resource exhaustion*, we exploit the naive assumptions in the SPS algorithm about other vehicles' transmission patterns, as well as its perceptual shortcomings about spectrum resources' availability, to induce a perception that the spectrum is far more crowded than it is and, consequently, induce an increase in packet collisions. This attack exploits elements of SPS which are common to both LTE-V2X and NR-V2X, revealing vulnerabilities in both protocols. SPS is designed under the assumption that vehicles transmit BSMs with constant periodicity, but that is not required to comply with the specifications. Using our hybrid hardware testbed, we demonstrate how an attacker can make strategic, *legitimate* transmissions in such a manner that other vehicles will misperceive the amount of available resources and avoid using certain parts of the (available) spectrum. We then show that the attack's cascading effects against multiple vehicles result in a large number of vehicles competing for the same, limited amount of bandwidth, increasing the rate of collisions and causing a much greater amount of packet loss. Since this attack only makes C-V2X-compliant transmissions, it is both stealthy and deniable (in the unlikely event of detection) under the current specifications. We further put forward preliminary experimental evidence for a promising detection approach based on regression analysis of channel usage over time.

We also propose mitigation approaches to address the vulnerabilities we identify in the C-V2X PHY layer and the SPS algorithm. Specifically, we show that minor modifications to the periodicity of BSMs can effectively mitigate targeted sidelink jamming with minimal collateral effects on SPS performance. Then, we propose and evaluate delicate adjustments to the listening period of SPS which reduces the effectiveness of sidelink resource exhaustion by more than 50%, providing insight for a more secure development of Rel-18 NR-V2X.

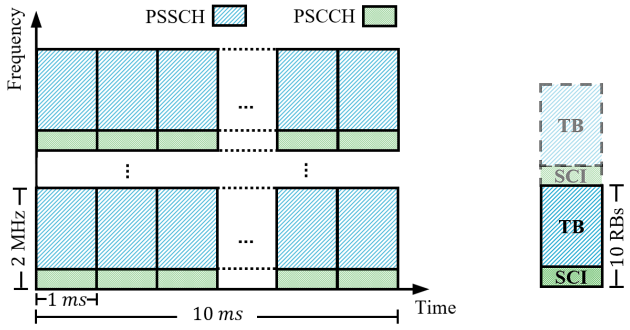

(a) Sidelink frame structure with 2 MHz subchannels. (b) A BSM.

Fig. 1: LTE-V2X sidelink frame structure and the TB and SCI elements of a message transmitted on at least one (usually two adjacent) subchannel(s). NR-V2X is identical in the time domain, but is not as rigidly defined in the frequency domain.

II. Technical Background

We begin with background on the PHY and MAC layers of C-V2X, including the design principles of the SPS algorithm.

A. Sidelink Communication

Sidelink communication allows direct communication between cellular devices synchronized using global navigation satellite systems (GNSS). Mode 4 (LTE) and Mode 2 (NR) are the only sidelink modes that allow such communication in situations where neither device is "in-coverage" (i.e., without requiring access to an eNB or gNB). As such scenarios are commonplace (e.g., on rural roads), we assume these modes throughout this paper. LTE and NR sidelink interfaces are both supported by 5G C-V2X systems [9], [23], as established by the most recent V2X specification (Rel-16 [3]) and cemented for long-term coexistence in the near-final draft of Rel-17 [10].

B. C-V2X PHY Layers

LTE-V2X and NR-V2X feature the same time-domain structure [24], with 10 ms sidelink frames divided into 1 ms subframes (see Fig. 1(a)). Each subframe is considered a time slot within which one or more transmissions can occur. Devices identify frames by sequential system frame numbers (SFNs) and subframes by sidelink frame index (SFI) between $0-9$.

In the frequency domain, LTE-V2X channels can be either 10 or 20 MHz wide [1]. NR-V2X channels are more flexible, allowing channels as wide as 400 MHz; however, typical configurations remain in the $10-20$ MHz range [25]. Without loss of generality, we consider only the 10 MHz configuration, which is common to LTE- and NR-V2X. A 10 MHz LTE-V2X channel is divided into five subchannels. The 3GPP specifications then allow two configurations for dividing each subchannel between control and shared (data) channels. Without loss of generality, we focus on the so-called "adjacent" configuration in which the channel is divided into contiguous 2 MHz subchannels, as shown in Fig. 1(a).

A subchannel in each subframe further consists of 10 LTE resource blocks (RBs). The first two RBs of each subchannel are used for the sidelink control channel (PSCCH) and

the remainder for the sidelink shared channel (PSSCH)–see Fig. 1(a). Any transmission within a subchannel then consists of sidelink control information (SCI), which is transmitted over the PSCCH, and a transport block (TB), which carries the payload in the PSSCH–see Fig. 1(b). As BSMs are sent every 100 ms [26], and assuming each BSM needs two subchannels (the common practice), a vehicle can choose one of the 400 subchannel pairs in a 100 ms period. Note that irrespective of the number of subchannels used, any single transmission must occur entirely within one subframe. Note also that no TB (in PSSCH) can be decoded without first decoding the associated SCI message (from PSCCH) [18], [19]. This is a critical point for the attack we present in Section III-B.

C. Semi-persistent Scheduling Algorithm

In LTE sidelink Mode 4 and NR sidelink Mode 2, vehicles use an autonomous MAC-layer protocol, *semi-persistent scheduling* (SPS), to choose the subframe and subchannel(s) they will use to transmit their periodic messages (e.g., BSMs) [18], [19]. Excluding minor differences, SPS is essentially the same in LTE- and NR-V2X [24]. As a sensing-based protocol, it is intended to allow short-term prediction of future channel usage based on a brief sensing (listening) period, allowing a vehicle to identify time and frequency resources not expected to be used by other vehicles. We note that C-V2X does not employ any power or multiple access control mechanisms (e.g., carrier sense before transmission), so SPS is the *only* means by which vehicles can try to avoid interfering with each others' transmissions. The scheduling in SPS is semi-persistent to allow vehicles to *reselect* resources every c messages so as to adapt to dynamic environments; for BSMs with the standard 10 Hz periodicity, $c \in \{5, \ldots, 15\}$ is randomly set every time an attempt is made. After c messages, the vehicle will decide with globally pre-configured probability $P \in \{0, 0.2, 0.4, 0.6, 0.8\}$ whether to select new resources (see below) or continue using its current selection.

SPS requires vehicles to constantly monitor channel usage and record the reference signal received power (RSRP) for each RB [18], [19]. When resource reselection is triggered, the last 1000 (in LTE-V2X) or 1100 (in NR-V2X) subframes' worth of sensing data are used to assess the set of candidate single-subframe resources (CSR) of a BSM period from which new resources will be selected (i.e., 500 resource options for 100 ms BSM period). The size of CSR is next reduced by filtering out any options that meet all of the following criteria during the listening period [18]:

1) At least one valid SCI message was received in PSCCH.
2) At least one valid TB was received in PSSCH using the resources indicated by the SCI message.
3) The average RSRP for the TBs of the resources with valid SCI and TB exceeds a given threshold, TH_{rx}.

If CSR is reduced to less than 20% of its original size, indicating a high noise/interference environment, the process is repeated with TH_{rx} increased by 3 dB. Otherwise, new resources are randomly chosen from CSR for subsequent periodic transmissions. Note that this process does not prevent vehicles from choosing conflicting resources. If two (or more) vehicles in the same area perform resource reselection at the same time, then their respective CSRs are likely to be similar and they may select the same resources; consequently, their packets will consistently collide until one or more of those vehicles performs another resource reselection and chooses different resources. The likelihood of expected packet loss due to this shortcoming increases with the number of vehicles, an important consideration for the stealthiness of DoS attacks which cause packet loss.

III. Proposed C-V2X DoS Attacks

In this section, we present novel, protocol-aware DoS attacks to exploit the shortcomings of C-V2X identified above.

A. Threat Model

We consider a single protocol-aware attacker, Eve, who has different disruptive goals in each attack. We generally assume she is capable of mimicking an ordinary V2V-equipped vehicle: she may be mobile or stationary and can communicate sidelink signals on, e.g., the 5.9 GHz band. We also assume Eve wishes to remain stealthy by avoiding detection (and its consequences); therefore, we require her to comply with all LTE-V2X specifications to appear outwardly legitimate:

- Eve may transmit at up to, but not beyond, the standard C-V2X power level of 23 dBm [27].
- Eve must synchronize with GNSS and transmit periodically at a valid C-V2X rate (20, 30, 50, or 100 ms) [24].
- Eve must comply with SPS requirements to regularly reselect resources [18].

B. Attack 1: Targeted Sidelink Jamming

In carrier-sense multiple access (CSMA) protocols like DSRC, it is nearly impossible to predict exactly when any particular vehicle will transmit a BSM, as medium contention (and related latency) precludes precise BSM periodicity. However, due to the rigid structure of C-V2X at the PHY-layer, observing the resources that a vehicle uses to transmit one BSM allows an observer (or attacker) to precisely predict the resources that vehicle will use for its next several BSMs. Through *targeted sidelink jamming*, we show how this can be leveraged by an attacker to jam the BSMs of a targeted vehicle. We assume Eve knows a specific victim vehicle (Alice) and wants to put it at increased risk of collision by preventing Alice's BSMs from being received by other vehicles.

1) Attack Procedure: In order to jam only Alice's BSMs, Eve must be able to first identify Alice's BSMs, which is potentially tricky due to the pseudonymization of BSM identifiers as per the IEEE 1609.2 V2V security standard [28]. However, we note that BSMs are never encrypted, and commercial standards (e.g., [29]) require BSMs to also contain potentially identifying information (e.g., color, make, model, length, width) about the sending vehicle. Further, techniques like angle-of-arrival estimation may be employed (e.g., if equipped with multiple antennas) to isolate Alice's BSMs from others.

Eve executes the following series of steps to attack Alice:

1) *Listen*: Eve continuously listens to the channel until she detects a BSM from Alice.
2) *Record*: Once such a BSM is received and processed, Eve marks the resources that Alice is currently using.
3) *Predict*: Eve further identifies in the sidelink resource grid the resources that Alice's next $c \in \{5, \ldots, 15\}$ BSMs will use.
4) *Jam*: In the predicted resources, Eve transmits a short SCI message to collide with Alice's, rendering the associated TB unrecoverable.
5) *Monitor*: Between jamming instances, Eve actively listens for possible BSMs from Alice in different resources.
6) *Update*: If monitoring uncovers that Alice has reselected resources, Eve goes back to step 2 to update her record and continue jamming Alice's BSMs.

The critical steps of the attack, *predict* and *jam*, are both facilitated by the PHY-layer design of C-V2X. We focus our discussion on LTE-V2X for clarity, although the same slot-based design is also used in NR-V2X. As BSMs are sent every 100 ms, Eve can determine from the SFN of Alice's first BSM, denoted by SFN_1, that future BSMs from Alice will arrive in frames SFN_{1+10i}, $i \in \{0, 1, 2, ...\}$ and the same SFI and subchannel(s) as the first BSM within those frames, until Alice performs another SPS resource reselection. Thus, Eve can anticipate and react to (i.e., jam) Alice's BSMs with extremely high accuracy. If Alice reselects new resources and her next BSM arrives earlier than expected, the *update* step ensures Eve will correct herself for the next frame; missing jamming no more than one BSM (in Section V, we show this has negligible impact on Eve's success).

With respect to the *jam* step, Eve can render an entire BSM irrecoverable by jamming only its associated SCI message in both LTE- and NR-V2X (see Section II-B). This allows Eve to knock out Alice's entire BSM (which in practice comprises at least 20 RBs) with little effort, by jamming only the 2-RB SCI message in PSCCH–a duty cycle of at most 10%.

2) Attack Detection: The primary BSM-DoS detection mechanism in V2V is based on monitoring packet delivery ratio (PDR) and reporting anomalously low values [30]–[33]. To assess the detectability of our attack, we investigate to what extent current PDR-based approaches are suitable.

C-V2X system model— To do this, we simulated a 10 MHz C-V2X channel in MATLAB assuming 5 subchannels, drawing on MATLAB's *LTE Toolbox* [34] for accurate C-V2X PHY-layer structure. We use a matrix M to represent subframes and subchannels, allowing us to track transmissions over simulated time and record packet collisions. In our model, any particular subchannel-subframe resource is denoted by $M_{i,j}$ where $i \in [1, 5], j \in \mathbb{Z}^+$. Following C-V2X standards, we configured simulated vehicles to transmit 2–subchannel-wide BSMs (per [35]) at the standard 10 Hz rate in this "channel." We implemented the SPS algorithm (per [18]) and configured the vehicles to perform SPS resource reselection at realistic intervals, allowing us to evaluate SPS packet loss against that caused by our attack.

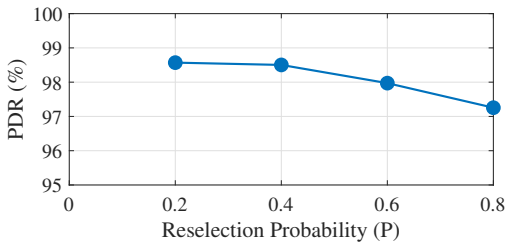

Fig. 2: PDR for different reselection probabilities P.

We deliberately chose not to simulate noise or other channel impairments in order to exclusively study packet loss due to SPS resource conflicts and DoS attacks. This creates a worst-case scenario for the attacker, who wishes to remain undetected, and a best-case scenario for detection, because Alice (or an independent monitor) will be able to statistically distinguish packet loss due to a DoS attack (from the packet loss which normally results from SPS) with fewer false alarms or missed detections. Thus, we assume every transmitted BSM will be successfully received and decoded by every vehicle unless it either collides with another vehicle's BSM (due to SPS resource conflicts) or is blocked by our DoS attack. We argue that if the effects of our attack cannot be reliably distinguished from benign packet loss under this model, then our attack will be even more challenging to detect in a channel where noise, fading, etc. contribute to additional packet loss.

Additionally, to find a resource reselection probability P that results in the fewest conflicts and so makes the attack more difficult to go undetected, we ran several simulations for each possible value of P, varying the number of vehicles between each iteration. Based on our results (shown in Fig. 2), we set $P = 0.2$ to minimize SPS-induced packet loss.

System monitor— We consider a system monitor who attempts to detect a DoS attack by monitoring overall PDR in the LTE-V2X channel. We assume the monitor can accurately estimate the number of vehicles in its area (e.g., based on historical traffic data [36]) and, therefrom, the number of BSMs it should receive in a given time period, for devising a statistical test threshold PDR_{th} below which an observed PDR value should trigger a DoS attack alert. Because SPS-induced packet loss increases with the number of vehicles (denoted by n_v) who are sharing the channel [37], PDR_{th} will always be a function of n_v, irrespective of the specific test to be used.

To derive the detection threshold, the monitor will calculate its observed PDR (PDR_{mon}) over an interval of t seconds as a test statistic based on n_v, the rate at which vehicles transmit BSMs (r), and the number of BSMs that the monitor successfully decodes in that interval (b). As r and t are known a priori, PDR_{mon} would be a function of n_v and b, as expressed by:

$$PDR_{\text{mon}}(n_v, b) = \frac{b}{n_v r t} \tag{1}$$

Now, we can detect the attack if

$$PDR_{\text{mon}}(n_v, b) < PDR_{\text{th}}(n_v) \tag{2}$$

Note that (2) can be used for detection irrespective of the specific definition of $PDR_{\text{th}}(n_v)$. Further, in our model, a number of packets are lost in packet collisions due to SPS (denoted by $b_{\text{SPS}}^{\text{lost}}$) and more packets are lost due to jamming attacks (denoted by $b_{\text{jammed}}^{\text{lost}}$). So, we can express b by

$$b = b_{\text{sent}} - b_{\text{SPS}}^{\text{lost}} - b_{\text{jammed}}^{\text{lost}} \tag{3}$$

Combining (1) through (3) yields the following

$$\frac{b_{\text{sent}} - b_{\text{SPS}}^{\text{lost}} - b_{\text{jammed}}^{\text{lost}}}{n_v r t} < PDR_{\text{th}}(n_v) \tag{4}$$

which helps to illustrate why PDR is not a reliable metric for DoS detection. Because SPS inevitably causes packet loss, particularly for larger n_v, $PDR_{\text{th}}(n_v)$ must *always* allow for some probabilistic range of packet losses (e.g., using standard error or a confidence interval) to avoid raising a false alarm. Since $PDR_{\text{th}}(n_v)$ is expected to incorporate these anticipated packet losses, if $b_{\text{jammed}}^{\text{lost}}$ were removed from the left-hand side of (4) then the inequality would almost never be true. As such, whether or not the attack is detectable depends in practice solely on the number of messages an attacker jams in an interval of t seconds. If an intelligent attacker ensures $b_{\text{jammed}}^{\text{lost}}$ is sufficiently small as to not satisfy (4), her DoS attack will be very difficult to detect.

We illustrate the difficulty of detecting targeted sidelink jamming using PDR through an example. Consider a monitor who defines $PDR_{\text{th}}(n_v)$ based on an estimate of PDR over time. In particular, we let $PDR_{\text{th}}(n_v)$ be the least-squares regression line for the lower bound of a 95% confidence interval on the average PDR, as calculated over 1-minute intervals. We ran simulations of the channel over intervals of 1 minute for all practical values of $n_v \in \{0, \ldots, 200\}$[1], calculating the PDR on a per sidelink frame basis. Based on 6000 PDR measurements for each n_v, we calculated the mean PDR and the 95% confidence interval, recording the lower bound of each interval. From those lower bounds, we calculated a least-squares regression line to represent $PDR_{\text{th}}(n_v)$:

$$PDR_{\text{th}}(n_v) = -0.0002 n_v + 1.0027 \tag{5}$$

One may infer from (5) that $PDR_{\text{th}}(n_v)$ will decrease and the confidence interval will widen as n_v increases, making any DoS effect inherently less detectable as n_v increases. We consider an attacker who varies the length of her attack within t seconds, and hence, $b_{\text{jammed}}^{\text{lost}}$ in (4), and we evaluate whether the attack can be reliably detected. As shown in Fig. 3, attacks of duration 15 or 30 seconds become statistically indistinguishable from benign packet loss at higher levels of n_v; further, the curves for every attack duration fall within 1% of SPS packet loss for all $n_v > 100$. Based on these results, which illustrate that our attack is difficult to detect using PDR in even a perfect channel, we argue a better detection approach is required.

[1]With 5 subchannels per subframe and BSMs spanning two subchannels, at most two vehicles can transmit within the same subframe.

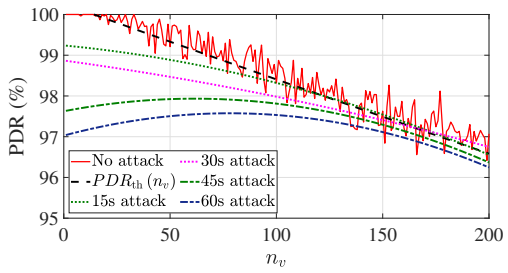

Fig. 3: Overall PDR in the LTE-V2X channel during targeted sidelink jamming under varying attack durations for all practical values of n_v and $t = 60$ s.

Attack detection through cluster analysis— Although our attack aims to have little impact on overall PDR, it does have a substantial impact on PDR for the specific frames wherein messages are being jammed. While SPS packet loss occurs across all frames, packet loss from our attack occurs only within certain frames (used by the target). Thus, by looking at PDR by frame index ($SFN \bmod r$), we posit that detection is possible by using cluster analysis to differentiate frames with benign packet loss due to SPS from those where packet loss is due to SPS and the attack combined.

This approach requires monitoring the channel for short periods (e.g., 10 seconds) and recording how many messages are received in each frame versus how many are expected. The monitor can then group together per-frame PDR measurements by frame index and compare PDR values for each index to check for anomalies using a cluster analysis algorithm like DBSCAN [38]. DBSCAN is an unsupervised algorithm, desirable for highly dynamic environments like C-V2X, and its worst-case $\mathcal{O}(\log(\text{n}))$ time complexity (see [38]) is suitable for execution on resource-constrained vehicular chipsets or standalone roadside infrastructure units. In Section V-A3, we show DBSCAN can in fact effectively cluster PDR values by frame index to identify outliers and, subsequently, the frame with significantly more outliers than others, an indication of a possible attack.

C. Attack 2: Sidelink Resource Exhaustion

In both LTE- and NR-V2X, SPS is used to predict future channel usage based on observations made during a short listening period. SPS works well as long as vehicles stick to transmitting periodic BSMs at consistent intervals. However, if an attacker breaks from this pattern while still complying with all C-V2X specifications, then the entire system can cripple. In this attack, we assume Eve wants to cause a DoS effect in the C-V2X channel without *directly* jamming messages from other vehicles. Through the *sidelink resource exhaustion attack*, we show how she can violate the spirit, but not the letter, of C-V2X specifications, and in doing so reduce C-V2X channel capacity and correspondingly increase packet loss for a large number of vehicles.

1) SPS Vulnerability: When a vehicle performs SPS resource reselection, its objective is to select new time-frequency

resources to use for transmitting its BSMs. During the listening period, vehicles observe the patterns of periodic transmissions from other vehicles and attempt to predict which resources will (and will not) be used in future frames. Normally, this works well because vehicles stick to the same periodicity over time, so a vehicle which sends a BSM in some specific time slot across the previous several frames can be expected to continue this pattern for the immediate future. However, since SPS assumes all transmissions will follow this pattern, it is vulnerable to an attacker whose transmissions are not actually of a given periodicity. For example, during the listening period, an attacker may transmit several times in a particular resource, then switch to another, and still a third, all within the time interval that another vehicle is performing SPS listening. This vehicle will not know that one attacker has made several transmissions across different resources and will assume three (or more) vehicles are actually using, and will continue to use, those resources. Thus, the vehicle will avoid selecting those resources. On a larger scale, this means a single attacker can cause many vehicles to avoid using part of the available channel and compete for a narrower bandwidth, which inevitably leads to vehicles more frequently selecting conflicting resources and increasing the rate of packet collisions.

This is all possible because, in both LTE- and NR-V2X, the standard SPS listening period is at least $1000\,\text{ms}$ long, ten times larger than the BSM period for which a candidate single-subframe resource set CSR is created for choosing new resources (see Section II-C). Therefore *the value of any particular radio resource in CSR is based not on one, but on several (at least 10) different radio resources* observed during the listening period. SPS's lack of perceptual granularity (i.e., the dependence of each candidate resource in CSR on more than one resource in the listening period) constitutes an exploitable vulnerability in the MAC layer of both C-V2X protocols.

2) Attack Design: Behind the scenes, Eve is strategic about the size and periodicity of her transmissions as well as the choices she makes during SPS. We require her to regularly reselect resources, and further let her deliberately select her new resources. For example, if Eve wants to prevent other vehicles from using a certain subchannel, she may anonymously transmit in that subchannel very frequently (e.g., every $20\,\text{ms}$). In combination with other vehicles that are legitimately using that subchannel, it will appear to SPS that the subchannel in question is completely in use, so no vehicles will attempt to use it. This holds true even if Eve only transmits in this manner for a fraction of a second and then switches to using different resources, which in turn compounds the effects. In Section V-B, we show how Eve can use this approach to push other vehicles away from using so much bandwidth that packet loss increases by up to 50% as vehicles compete for what few resources they believe are still available.

3) Attack Detection: Due to Eve's strict compliance with C-V2X specifications, detecting sidelink resource exhaustion based on PDR is tricky. A monitor could easily observe a 50% decrease in channel PDR; however, the cause will not be evident. As Eve does not directly jam BSMs, there is no observable alignment of her transmissions with lost packets. Further, neither the size of her transmitted BSMs, her rate of BSM transmission, nor her transmit power deviate from those of ordinary vehicles. Thus, we contend that attempting direct detection of Eve (i.e., attempting to identify her malicious transmissions) is unlikely to succeed.

Instead, we propose that a monitor should look at channel resource usage levels over time for the specific effects of sidelink resource exhaustion in order to infer Eve's presence. This attack has the particular effect of causing vehicles to use only part of the available channel; i.e., a large percentage of channel bandwidth will be wasted. This narrowing effect does not occur under normal SPS operation and can be considered as a sort of signature for the sidelink resource exhaustion attack. Then, by monitoring channel resource usage, it is possible to observe when the number of used resources diminishes unexpectedly, facilitating detection of the attack. One way to monitor channel resource usage in this fashion is to approximate trends using least-squares regression analysis. We demonstrate the effectiveness of this approach in Section V-B4.

IV. C-V2X DoS Attack Mitigation

In this section, we put forward steps to address and mitigate each of our DoS attacks against C-V2X.

A. Targeted Sidelink Jamming

The targeted sidelink jamming attack works by exploiting the precise BSM periodicity that results from the slot-based LTE- and NR-V2X PHY layers and GNSS synchronization. Thus, inducing variation in message periodicity is a potential mitigation technique. Due to the use of past observations in SPS to predict channel usage, periodicity cannot be completely eliminated; however, we propose moving periodicity to a per-vehicle rather than global definition. Specifically, we propose requiring vehicles to slightly modify their periodicity each time they reselect resources, e.g., $c_1 \in \{5, \ldots, 15\}$ BSMs spaced $97\,\text{ms}$ apart might be followed by $c_2 \in \{5, \ldots, 15\}$ separated by $102\,\text{ms}$, and so on. This mitigation would have negligible impact on end-to-end BSM latency, and it would significantly reduce the ability of Eve to predict future BSMs based on just one. Eve would instead have to listen to at least two BSMs to identify her target's periodicity before her jamming could begin; further, her jamming may be more easily identified as tracking with a particular vehicle's BSM periodicities, increasing the detectability of the attack.

The key question is to what extent this mitigation would complicate SPS' identification of a *set* of resources to use and the impact of variable periods on its prediction performance, as the current SPS algorithm is predicated on a globally common BSM periodicity (e.g., $10\,\text{Hz}$ for all vehicles). We took steps to answer this question and discuss our results in Section V-C1.

B. Sidelink Resource Exhaustion

Our sidelink resource exhaustion attack exploits the disparity in size between the listening period and CSR, a

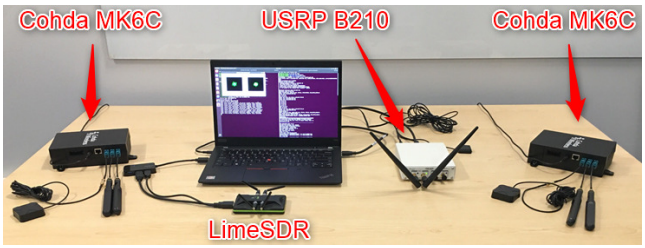

Fig. 4: Experimental testbed for targeted sidelink jamming.

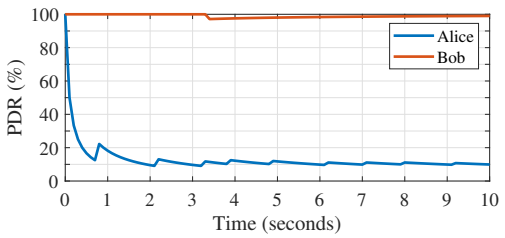

Fig. 5: Testbed results showing reduction in target PDR (up to 90%) due to targeted sidelink jamming.

fundamental vulnerability in SPS. Thus, one mitigation would be modifying SPS so that disparity no longer exists. This could proceed in two directions. First, the length of the SPS listening period could be reduced from 1000 ms to 100 ms (the size of the CSR). This would eliminate the influence of earlier, less relevant transmissions (including the attacker's); in fact, Rel-16 NR-V2X does this, but exclusively for aperiodic messages [39]. Whether this technique could be applied to periodic traffic is unknown and its impact will be investigated in our future work. Second, CSR might be enlarged to 1000 ms to match the size of the SPS listening period. This is similar to some proposals for SPS performance enhancement (e.g., [20], [21], [40]), which propose extending the interval between resource reselections to reduce the chance that vehicles choose conflicting resources by reducing the overall number of reselections. Here, rather than reducing the number of reselections, we propose to give a similar level of flexibility by allowing vehicles a greater choice of resources to use during each reselection. We thereby deny Eve the opportunity to improperly influence resource selections, as her rate-limited transmissions would have less impact on a larger candidate resource set than on the current configuration.

V. Performance Evaluation

We experimentally validated the effectiveness of our attacks against state-of-the-art commercial LTE-V2X equipment using our hybrid hardware testbed, consisting of Cohda MK6C V2X on-board unit (OBU) evaluation kits with Qualcomm 9150 chipset [22] and one or more Ettus USRP B210s (see Fig. 4). In all experiments, each OBU was equipped with two 4 dBi antennas and transmitted at a power level of 23 dBm (as is standard in LTE-V2X [27]) while our USRP(s), each equipped with two 5 dBi antennas, transmitted at a power level of approximately 15 dBm [41] within a Faraday lab at RIT's Global Cybersecurity Institute. The devices all used a GPSDO module for GNSS time synchronization, but in the absence of real GNSS signal in the Faraday lab, we used a LimeSDR and GPS-SDR-SIM [42] to generate synthesized GNSS signals. We implemented both of our attacks in C++ by extending srsRAN [43] and work by Eckermann and Wietfeld [44].

We further demonstrate the effectiveness of our detection techniques and mitigation approaches through MATLAB simulations. We make some careful assumptions (e.g., a perfect channel—see Section III-B), ensuring our approaches remain compatible with and extendable to a real V2V environment.

A. Targeted Sidelink Jamming

1) Testbed configuration: To represent two vehicles, referred to as Alice and "Bob", we configured two OBUs to transmit BSMs at the standard 10 Hz rate. Alice and Bob were placed 1 m apart while Eve was positioned 2 m from both Alice and Bob. Eve was represented by one USRP.

2) Experimental results: We evaluated the effectiveness of the attack based on Eve's ability to degrade Alice's PDR. We first ran Alice and Bob for 10 minutes without an attacker to obtain a baseline PDR measurement. Comparison of Bob's received packet log against Alice's transmission log showed Bob received >99.85% of the BSMs sent by Alice when no attacker was present. We then repeated this experiment, adding Eve. We again compared the packet logs; as shown in Fig. 5, our attack was very successful. Eve was able to reduce Alice's PDR by 90% in less than two seconds while leaving Bob's messages largely untouched. These results confirm the effectiveness, precision, and real-world viability of our attack against state-of-the-art LTE-V2X OBUs.

3) Detection with DBSCAN: To demonstrate how DBSCAN can be used to detect targeted sidelink jamming, we simulated the attack (based on our model from Section III-B2) and recorded per-frame PDR (averaged per 10 frames) over a 100-second period. Based on sorted k-Nearest distance, we set $\epsilon = 0.1$ as the geometric distance parameter for DBSCAN to identify clusters. The attacker in the simulations targeted a vehicle which generally transmitted its BSMs in the seventh of every ten frames. As Fig. 6(b) depicts, DBSCAN identified more than twice as many anomalously low PDR measurements for frame index 7 than for other frame indices, establishing the effectiveness of this technique. The original (unoptimized) DBSCAN algorithm, while effective, may be too slow in practice, but our results suggest an optimized variant (e.g., OPTICS [45], HDBSCAN [46]) may be both effective and practical. We leave further investigation to future work.

B. Sidelink Resource Exhaustion

1) Testbed configuration: To verify our premise that strategic transmissions by Eve can influence the SPS resource selections of other vehicles, we set up an experiment involving one OBU and two USRPs. As before, we configured the OBU to transmit BSMs at the standard rate of 10 Hz and we used one USRP for Eve, configured to transmit BSMs every 20 ms (i.e., in every other sidelink frame) using the first two subchannels of the first sidelink subframe. We used another

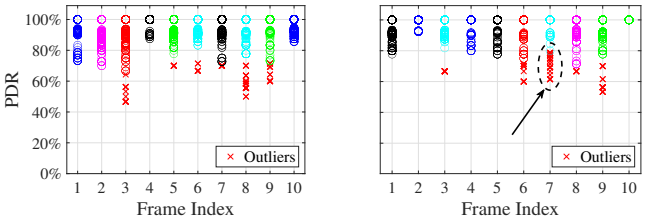

(a) Without attack. (b) With attack.

Fig. 6: DBSCAN clustering results to analyze PDR. In (b), the effects of the attack are marked with an ellipse and arrow.

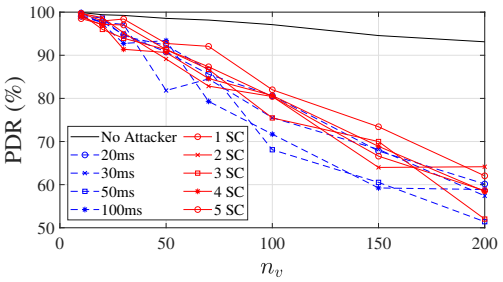

Fig. 7: Impact of sidelink resource exhaustion on channel PDR versus the number of vehicles (n_v) for different attacker transmission size (in subchannels (SC)) and periodicities.

USRP to monitor and record the PHY-layer resources chosen by the OBU running SPS.

2) Experimental results and analysis: We were able to confirm our premise, finding that over a period of 116 minutes wherein the OBU performed at least 4500 resource reselections, it *never* selected the resources used by Eve—not only in the frames Eve transmitted in, but in *all* frames–even while selecting every other available resource at least several times over the same period. A 116-minute experiment is necessary in order to rule out the possibility that the OBU does not select Eve's resources solely due to probability, as explained below.

Given a 10 MHz channel with 5 subchannels, there are 4 pairs of adjacent subchannels from which one can be selected to carry a 2-subchannel BSM, and thus there are $100 \times 4 = 400$ total candidate resources per BSM interval. The chance of the OBU selecting any particular resource is therefore $\frac{1}{400} = 0.0025$, and the chance of selecting one or both of the two resources Eve is transmitting over is

$$\mathrm{P}(\mathrm{R1} \vee \mathrm{R2}) = 0.0025 + 0.0025 = 0.005$$

where R1 and R2 are subchannel pairs $1{-}2$ and $2{-}3$, respectively. Thus, the chance of a device *not* selecting either of those resources is $\neg\mathrm{P}(\mathrm{R1} \vee \mathrm{R2}) = 0.995$. Considering also that resources are only reselected with probability P, an unknown value for the closed-source Cohda OBUs, this probability is actually $0.995P$. We assume a worst-case value of $P = 0.2$, which corresponds to the fewest resource reselections over time, so the probability of the OBU not choosing Eve's resources in any given reselection is

$$\neg\mathrm{P}\left(\text{reselect}\right) \vee \left[\mathrm{P}\left(\text{reselect}\right) \wedge \neg\mathrm{P}\left(\mathrm{R1} \vee \mathrm{R2}\right)\right]$$
$$(1 - 0.2) + (0.2 \times 0.995) = 0.999$$

Over n reselections, then, the probability that a device never selects the resources used by Eve is 0.999^n. For this probability to be less than 1%, $n = \log_{0.999}(0.01) = 4603$ reselections would be necessary. Knowing a device performs resource reselection at most once every 1500 ms (i.e., after 15 transmissions), the experiment needed to be run for 4603×1500 ms (115.7 minutes) in order to claim, with 99% confidence, that the OBU's avoidance of Eve's resources was due to her transmissions rather than simple probability.

3) Simulations: Having confirmed through experiments that Eve can influence a vehicle's SPS selections, we now evaluate the effectiveness of our attack against a C-V2X channel used by varying numbers of vehicles through MATLAB simulations. We used the same channel model as in Section III, and note that while our model is an LTE-V2X channel, the modeled features (PHY structure and SPS) are essentially identical to the design of NR-V2X. Therefore, our evaluation of this attack applies to both technologies.

We also evaluate the effects of the attack on overall PDR based on the size (in subchannels) and the periodicity of Eve's transmissions. First, we held Eve's transmission periodicity constant at 20 ms and varied her transmission size between $1-5$ subchannels. Then, we held Eve's transmission size constant at 2 subchannels (the practical BSM size) and varied her transmission periodicity between the allowed values of 20, 30, 50, and 100 ms. Fig. 7 exhibits the results for all of these scenarios, showing that in all cases, Eve's actions result in significantly more packet loss than SPS alone, in some cases as much as 50% more. Fig. 7 also shows our attack is more effective when a greater number of vehicles are using the channel. This is a direct consequence of our attack design: as more vehicles are pushed by Eve into using fewer resources, it becomes more likely that conflicting resources will be selected.

We make two notes on this point. First, our results come from a perfect channel model where SPS conflicts are the only cause for packet loss. Therefore, what we have shown is actually the *minimum* amount of additional packet loss that Eve's actions will add on top of any packet loss that occurs from noise, interference, etc. in a real channel. Second, because our results show Eve's level of impact is directly related to the number of vehicles using the channel, sidelink resource exhaustion essentially flips SPS on its head. When the channel is busy, SPS is supposed to balance channel load across all available bandwidth, ensuring high throughput even under maximum load. However, we have demonstrated that our attack reduces C-V2X channel throughput under exactly these conditions, allowing an attacker to cause the most damage to C-V2X system performance in situations where it has the greatest need to perform at peak capacity.

4) Detection through regression analysis: To evaluate the effectiveness of our proposed detection approach, we ran

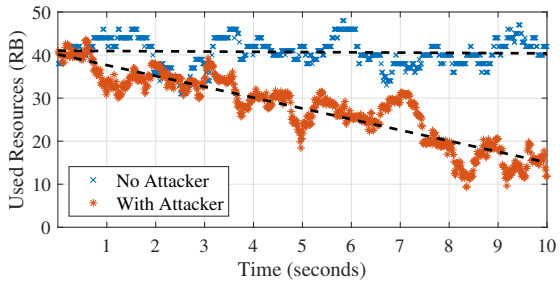

Fig. 8: Channel usage levels over time during the attack versus under normal operation.

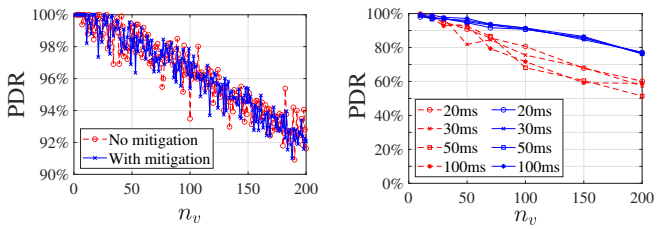

(a) Targeted Sidelink Jamming. (b) Sidelink Resource Exhaustion.

Fig. 9: Effectiveness of attack mitigations. Dashed (red) lines are without mitigation, solid (blue) lines are with mitigation.

the same simulations as above while monitoring per-frame resource usage over time. Then, we applied least-squares regression analysis to the collected data. As Fig. 8 shows, there is an obvious divergence of trends in resource usage when an attacker is or is not active. With no attacker, the slope of the regression line is $\delta = 2.4e{-}3$; as expected, this is a negligible change. When an attacker is present, though, regression analysis shows a strong negative trend of $\delta = -0.2$. More in-depth analysis and testbed experiments are required to establish a threshold for when the trend in resource usage is sufficiently negative to deem indicative of an attack; however, these results demonstrate the general effectiveness of this technique for detecting a sidelink resource exhaustion attack.

C. Mitigation Techniques

1) Targeted sidelink jamming: We now evaluate the effectiveness of allowing variable BSM periodicity as a mitigation proposed in Section IV-A. To do this, we simulated 10 minutes of our LTE-V2X environment and compared the amount of SPS-induced packet loss that occurred with and without the mitigation applied. As shown in Fig. 9(a), we observed no significant difference in PDR when vehicles are allowed to vary their periodicity between $90-110$ ms, indicating that adopting this mitigation approach is possible without causing a noticeable decrease in system performance.

2) Sidelink resource exhaustion: We evaluated the effectiveness of shortening the SPS listening period from 1000 to 100 ms, aligning it with the size of CSR. We then repeated the simulations from Section V-B, varying Eve's attack periodicity and BSM size, but this time with the mitigation applied. We found that, as shown in Fig. 9(b), shortening the SPS listening period substantially reduces the effectiveness of Eve's attack; while she still has a noticeable impact, she is not able to reduce PDR much below 80%, a significant improvement on the nearly 50% reduction that Eve can achieve without a mitigation in place. These results demonstrate the promising nature of our approach and support further refinement of the technique as a part of future work.

VI. Related Work

To the best of our knowledge, only two prior works describe protocol-aware C-V2X DoS attacks. Trkulja *et al.* [47] presented an attack where several colluding attackers collaboratively use SPS to predict and jam BSMs from other vehicles. The scope of this attack is limited as it requires multiple attackers; our attacks require just one. Also, whereas the system model in [47] does not match C-V2X standards; ours accounts for significantly more details of the C-V2X PHY/MAC layers. Li *et al.* [48] proposed a resource exhaustion attack against LTE-V2X based on flooding a network with high-priority packets. This attack targeted eNBs in LTE sidelink Mode 3, a significant difference from our work, which assumes sidelink Mode 4 and targets vehicles directly. Further, unlike Li *et al.*, our attacks cannot be mitigated by network-layer filters, and while their attack requires significant deviation from LTE-V2X standards, both of our attacks are protocol-compliant, significantly reducing their detectability.

BSM DoS attack detection is often based on monitoring PDR [31]–[33]. Unfortunately, such an approach is usually based on the assumption that DSRC will be the underlying V2V protocol. This assumption requires another, often unstated assumption that packet loss in the absence of an attacker (excluding environmental factors) will be negligible, due to the effective (if inefficient) medium contention mechanism used in DSRC. The use of SPS in C-V2X protocols invalidates this assumption, making detection techniques designed for DSRC generally inapplicable to C-V2X.

VII. Conclusion

C-V2X promises to help realize the safety benefits of V2V, but its effectiveness will be severely limited if security considerations are not promptly addressed. In this paper, we exposed fundamental vulnerabilities in the PHY and MAC layers of LTE-V2X and NR-V2X, the two protocols that will comprise 5G C-V2X systems of the near future. We devised, experimentally validated, and assessed the severity of two novel DoS attacks specifically engineered to exploit the shortcomings of C-V2X's slot-based PHY layer and SPS scheduling algorithm. We demonstrated the difficulty of detecting both attacks under current paradigms, proposed and experimentally validated promising new detection techniques, and further proposed and preliminarily evaluated mitigations for both attacks, thus providing direction to improve the security of C-V2X against protocol-aware DoS attacks.

REFERENCES

[1] Release 14 Description; Summary of Rel-14 Work Items, 3GPP Technical Report 21.914, Jun. 2018.

[2] Release 15 Description; Summary of Rel-15 Work Items, 3GPP Technical Report 21.915, Oct. 2019.

[3] Release 16 Description; Summary of Rel-16 Work Items, 3GPP Technical Report 21.916, Jun. 2021.

[4] Release 17 Description; Summary of Rel-17 Work Items, 3GPP Technical Report (Draft) 21.917, Sep. 2021.

[5] National Highway Traffic Safety Administration (NHTSA), "Notice of Proposed Rulemaking: Federal Motor Vehicle Safety Standards; V2V Communications," Federal Register, Vol. 82, No. 8, Jan. 2017, Accessed: Jul. 30, 2021.

[6] R. Wu, "C-V2X automotive tech brings enhanced safety and efficiency to China's roads," Mar. 2021, Accessed: Jun. 9, 2021. [Online]. Available: https://www.qualcomm.com/news/onq/2021/03/02/c-v2x-brings-enhanced-safety-and-efficiency-chinas-roads

[7] D. Butler, "How 'talking' and 'listening' vehicles could make roads safer, cities better," Jan. 2019, Accessed: Jun. 9, 2021. [Online]. Available: https://bit.ly/3zTZSYm

[8] 5G Automotive Association (5GAA), "A visionary roadmap for advanced driving use cases, connectivity technologies, and radio spectrum needs," Sep. 2020, Accessed: Jun. 8, 2021. [Online]. Available: https://5gaa.org/news/the-new-c-v2x-roadmap-for-automotive-connectivity/

[9] S. Patil, "How NR-based sidelink expands 5G C-V2X to support new advanced use cases," Qualcomm Technologies, Inc., Mar. 2020, Accessed: May 20, 2020. [Online]. Available: https://www.qualcomm.com/invention/5g/cellular-v2x

[10] Architecture enhancements for 5G System (5GS) to support Vehicle-to-Everything (V2X) services (Release 17), 3GPP Technical Specification 23.287, 2021.

[11] M. Sun, A. Al-Hashimi, M. Li, and R. Gerdes, "Impacts of constrained sensing and communication based attacks on vehicular platoons," *IEEE Trans. Veh. Technol.*, vol. 69, no. 5, pp. 4773–4787, 2020.

[12] H. Pirayesh, P. K. Sangdeh, S. Zhang, Q. Yan, and H. Zeng, "JammingBird: Jamming-resilient communications for vehicular ad hoc networks," in *Proc. IEEE Int. Conf. on Sensing, Commun. and Netw. (SECON)*, Virtual Conference, Jul. 2021.

[13] Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Wireless Access in Vehicular Environments, IEEE Standard 802.11p, 2010.

[14] R. Lu, L. Zhang, J. Ni, and Y. Fang, "5G vehicle-to-everything services: Gearing up for security and privacy," *Proc. IEEE*, vol. 108, no. 2, pp. 373–389, Feb. 2020.

[15] C. Lai, R. Lu, D. Zheng, and X. Shen, "Security and privacy challenges in 5G-enabled vehicular networks," *IEEE Netw.*, vol. 34, no. 2, pp. 37–45, Mar. 2020.

[16] Y. Liu *et al.*, "Secrecy rate maximization via radio resource allocation in cellular underlaying V2V communications," *IEEE Trans. Veh. Technol.*, vol. 69, pp. 7281–7294, Apr. 2020.

[17] X. Luo, Y. Liu, H.-H. Chen, and Q. Guo, "Physical layer security in intelligently connected vehicle networks," *IEEE Netw.*, vol. 34, no. 5, pp. 232–239, Sep. 2020.

[18] Physical layer procedures, 3GPP Technical Specification 36.213, 2021.

[19] Physical layer procedures for control (Release 16), 3GPP Technical Report 38.213, 2021.

[20] M. Chen, R. Chai, H. Hu, W. Jiang, and L. He, "Performance evaluation of C-V2X mode 4 communications," in *Proc. IEEE Wireless Commun. and Netw. Conf. (WCNC)*, Nanjing, China, Mar. 2021.

[21] A. Nabil, K. Kaur, C. Dietrich, and V. Marojevic, "Performance analysis of sensing-based semi-persistent scheduling in C-V2X networks," in *Proc. IEEE Veh. Technol. Conf. (VTC)*, Chicago, IL, USA, Aug. 2018.

[22] Cohda Wireless, "MK6c EVK - Cohda Wireless," 2020, Accessed: Oct. 26, 2020. [Online]. Available: https://bit.ly/2TCgCQt

[23] M. Demler, "C-V2X drives intelligent transportation," 2020, Accessed: Jul. 31, 2021. [Online]. Available: https://bit.ly/3txUSIK

[24] M. H. C. Garcia *et al.*, "A tutorial on 5G NR V2X communications," *IEEE Commun. Surveys Tuts.*, vol. 23, no. 3, pp. 1972 – 2026, 2021.

[25] W. Anwar, N. Franchi, and G. Fettweis, "Physical layer evaluation of V2X communications technologies: 5G NR-V2X, LTE-V2X, IEEE 802.11bd, and IEEE 802.11p," in *Proc. IEEE Veh. Technol. Conf. (VTC)*, Honolulu, HI, USA, Sep. 2019.

[26] G. Twardokus and H. Rahbari, "Evaluating V2V security on an SDR testbed," in *Proc. IEEE Conf. Computer Commun. Workshops (INFOCOM WKSHPS)*, (Virtual) Vancouver, BC, Canada, May 2021.

[27] Vehicle to Vehicle (V2V) services based on LTE sidelink; User Equipment (UE) radio transmission and reception, 3GPP Technical Report 36.785, 2016.

[28] Wireless Access in Vehicular Environments–Security Services for Applications and Management Messages, IEEE Standard 1609.2-2016, 2016.

[29] V2X Communications Message Set Dictionary, SAE International Standard J2735E, 2020.

[30] A. Benslimane and H. Nguyen-Minh, "Jamming attack model and detection method for beacons under multichannel operation in vehicular networks," *IEEE Trans. Veh. Technol.*, vol. 66, no. 7, pp. 6475–6488, Jul. 2017.

[31] N. Lyamin, D. Kleyko, Q. Delooz, and A. Vinel, "Real-time jamming DoS detection in safety-critical V2V C-ITS using data mining," *IEEE Commun. Lett.*, vol. 23, no. 3, pp. 442–445, Mar. 2019.

[32] ——, "AI-Based malicious network traffic detection in VANETs," *IEEE Netw.*, vol. 32, no. 6, pp. 15–21, Nov. 2018.

[33] A. Alipour-Fanid, M. Dabaghchian, H. Zhang, and K. Zeng, "String stability analysis of cooperative adaptive cruise control under jamming attacks," in *Proc. IEEE Int. Symp. High Assurance Syst. Eng.*, Singapore, Jan. 2017.

[34] The MathWorks, Inc., "LTE Toolbox," 2021, Accessed: Aug. 1, 2021. [Online]. Available: https://www.mathworks.com/products/lte.html

[35] Dedicated Short Range Communication (DSRC) Systems Engineering Process Guidance for SAE J2945/X Documents and Common Design Concepts, SAE International Standard J2945, 2017.

[36] S. Dongre and H. Rahbari, "Message sieving to mitigate smart gridlock attacks in V2V," in *Proc. ACM Conf. Secur. and Privacy in Wireless & Mobile Netw. (WiSec)*, Virtual Conference, Jul. 2021.

[37] 5GAA, "Test procedures and results for 20-MHz deployment in CH183," Tech. Rep., 2019, Accessed: Jul. 21, 2021. [Online]. Available: https://bit.ly/2V8vF8z

[38] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proc. Int. Conf. on Knowledge Discovery and Data Mining (KDD)*, Portland, OR, USA, Aug. 1996, p. 226–231.

[39] Overall description of Radio Access Network (RAN) aspects for Vehicle-to-everything (V2X) based on LTE and NR (Release 16), 3GPP Technical Report 37.985, Jul. 2020.

[40] S. Bartoletti, B. M. Masini, V. Martinez, I. Sarris, and A. Bazzi, "Impact of the generation interval on the performance of sidelink C-V2X autonomous mode," *IEEE Access*, vol. 9, pp. 35 121 – 35 135, Feb. 2021.

[41] M. W. O'Brien, J. S. Harris, O. Popescu, and D. C. Popescu, "An experimental study of the transmit power for a USRP software-defined radio," in *Proc. Int. Conf. Commun. (COMM)*, Bucharest, Romania, Jun. 2018, pp. 377–380.

[42] T. Ebinuma, "GPS-SDR-SIM," 2018, Accessed: Apr. 20, 2020. [Online]. Available: https://github.com/osqzss/gps-sdr-sim

[43] I. Gomez-Miguelez, A. Garcia-Saavedra, P. Sutton, P. Serrano, C. Cano, and D. Leith, "SrsLTE: An open-source platform for LTE evolution and experimentation," in *Proc. Tenth ACM Int. Workshop Wireless Netw. Testbeds, Experimental Eval., and Characterization (WINTECH)*, New York City, NY, USA, Oct. 2016, pp. 25–32.

[44] F. Eckermann and C. Wietfeld, "SDR-based open-source C-V2X traffic generator for stress testing vehicular communication," in *Proc. IEEE Veh. Technol. Conf. (VTC)*, Helsinki, Finland, Apr. 2021.

[45] M. Ankerst, M. M. Breunig, H.-P. Kriegel, and J. Sander, "OPTICS: ordering points to identify the clustering structure," *SIGMOD Record*, vol. 28, no. 2, p. 49–60, Jun. 1999.

[46] R. J. Campello, D. Moulavi, , and J. Sander, "Density-based clustering based on hierarchical density estimates," in *Proc. Pacific-Asia Conf. on Knowledge Discovery and Data Mining (PAKDD)*, Gold Coast, Australia, Apr. 2013, pp. 160–172.

[47] N. Trkulja, D. Starobinski, and R. A. Berry, "Denial-of-service attacks on C-V2X networks," in *Proc. ACM Workshop Automot. and Auton. Vehicle Secur. (AutoSec)*, Virtual, Feb. 2021.

[48] Y. Li, R. Hou, K.-S. Lui, and H. Li, "An MEC-based DoS attack detection mechanism for C-V2X networks," in *Proc. IEEE Global Commun. Conf. (GLOBECOM'18)*, Abu Dhabi, United Arab Emirates, Dec. 2018.