Towards Protecting 5G Sidelink Scheduling in C-V2X Against Intelligent DoS Attacks

Geoff Twardokus, *Graduate Student Member, IEEE,* and Hanif Rahbari, *Member, IEEE*

***Abstract*—5G Cellular Vehicle-to-Everything (5G C-V2X) is emerging as the globally dominant connected vehicle technology. One critical application of 5G C-V2X is the direct exchange of safety-critical messages between vehicles to prevent crashes and correspondingly reduce roadway injuries and fatalities. While current C-V2X security protocols concern only message payloads, we expose vulnerabilities in the physical-layer attributes and decentralized MAC-layer scheduling algorithm of 5G C-V2X by developing two stealthy denial-of-service (DoS) attacks to exploit them. These low-duty-cycle attacks dramatically degrade C-V2X availability, increasing the likelihood of prolonged travel times and even vehicle crashes. We further develop detection and mitigation techniques for each attack, in part by exploiting new C-V2X features of 3GPP Rel-17. We experimentally evaluate our attacks and countermeasures in a hardware testbed composed of USRPs and state-of-the-art C-V2X kits as well as through extensive network and roadway simulations, showing that within seconds of initiation our attacks can reduce a target's packet delivery ratio by 90% or that of the C-V2X channel to under 25%. We further evaluate our machine-learning detection and low-cost mitigation techniques, showing the latter completely thwart one attack and reduce the impact of the other by 80%, providing insight towards developing a more robust 5G C-V2X.**

***Index Terms*—V2V security, denial-of-service, selective jamming, resource scheduling, 5G**

I. Introduction

THE emergence of 5G Cellular Vehicle-to-Everything (C-V2X) is taking the intelligent transportation paradigm to the next level. A hybrid of the LTE-V2X and more recent New Radio (NR)-V2X communication protocols, 5G C-V2X has become the globally dominant suite of connected vehicle technologies with a nationwide deployment underway in China [1], exclusive rights to use the 5.9 GHz Intelligent Transportation Systems band in the U.S. [2], and increasing regulatory acceptance in the previously hesitant E.U. [3]. Connected vehicles employing C-V2X for vehicle-to-vehicle (V2V) communication take advantage of *sidelink* signaling (in contrast to typical uplink or downlink) to proactively broadcast periodic *basic safety messages (BSMs)* that allow vehicles to independently maintain awareness of each other's movements. The resulting awareness will facilitate collision avoidance and more efficient routing, particularly when neither a driver nor the onboard sensors of an advanced driver-assistance system can perceive an imminent collision; e.g., LiDAR sensors in non-line-of-sight (NLOS) scenarios. Widespread deployment and use of V2V in the U.S. is ultimately expected to prevent up to 600,000 vehicle crashes every year [4].

As a safety-critical technology, V2V must be extremely resilient to attacks. A vehicle moving at high speed may only have seconds of reaction time to avoid a collision; therefore, ensuring the *availability* of C-V2X at the physical (PHY) and MAC layers to support delivering BSMs is imperative. However, recent work on C-V2X has focused either on confidentiality or integrity (rather than availability) at the PHY and MAC layers (e.g., [5], [6]) or on security at the network or application layers (e.g., [7]–[9]). As far as we are aware, no prior work has thoroughly investigated *availability* vulnerabilities at the PHY layer of C-V2X, with only [10] investigating and exposing such a vulnerability at the MAC layer. However, the analysis in [10] is limited only to LTE-V2X, based on a simplified abstraction of its PHY layer and somewhat unrealistic assumptions about the MAC layer (e.g., assuming parameter choices that do not exist in the 3GPP standards). Therefore, a thorough examination of both NR-V2X and LTE-V2X under a realistic setup in search of both PHY- and MAC-layer vulnerabilities to threats against 5G C-V2X availability is essential.

In our preliminary work [11], we initiated such a thorough, more realistic investigation focused on LTE-V2X and exposed a wider set of PHY- *and* MAC-layer vulnerabilities by demonstrating novel, stealthy denial-of-service (DoS) attacks that exploit them. We further proposed preliminary solutions with promising results for detecting and mitigating those attacks. In this paper, we go beyond and in some aspects even overhaul that preliminary effort. We first extend our analysis to 5G C-V2X and demonstrate that the vulnerabilities we previously exposed in LTE-V2X persist in NR-V2X. More importantly, we refine and formalize our DoS attacks, presenting a new variant of one attack to inflict twice as much damage as its original incarnation in [11] (shown in this paper using a 3GPP-compliant 5G C-V2X simulator), and further study its impact on roadway traffic dynamics (e.g., travel time). We also present enhanced techniques for detecting each attack, with more extensive experiments and more effective mitigations than in [11], replacing one of our initial ideas from [11] with an entirely new approach based on new features of NR-V2X introduced in 3GPP Release 17.

Physical and MAC Layer Vulnerabilities: In 5G C-V2X, sidelink transmissions at the PHY layer must align with the time-frequency structure shown in Fig. 1(a), which is beneficial for, e.g., allowing simultaneous BSM transmissions and better coordinating channel resource selections. However, it also makes certain events highly predictable; in particular, this makes it easy to accurately anticipate when periodic

Geoff Twardokus is with the Electrical and Computer Engineering Ph.D. program and ESL Global Cybersecurity Institute at Rochester Institute of Technology, Rochester, NY, USA.

Hanif Rahbari is with the Electrical and Computer Engineering Ph.D. program and ESL Global Cybersecurity Institute at Rochester Institute of Technology, Rochester, NY, USA.

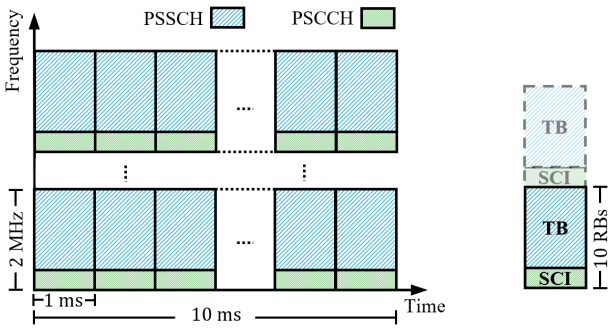

(a) Sidelink frame structure with 2 MHz subchannels. (b) A BSM.

Fig. 1: 5G C-V2X sidelink frame structure. The Sidelink Control Information (SCI) and Transport Block (TB) elements for a BSM are transmitted using 1 or more adjacent subchannel(s).

transmissions by particular vehicle(s) will occur. For BSMs, this structure results in a very precise periodicity, allowing an attacker to easily predict (and, ultimately, obstruct) transmissions with microsecond precision, as we show in this paper.

At the MAC layer, channel resource selection is decentralized, a requirement for C-V2X sidelink to avoid relying on base stations (gNBs in 5G) for resource allocation. The decentralized MAC-layer scheduling protocol used in 5G C-V2X is called *semi-persistent scheduling (SPS)* [12], [13]. As a sensing-based algorithm, SPS uses observations of historical patterns to predict short-term, and therefore semi-persistent, future channel usage. This allows vehicles to identify candidate time-frequency resources for their transmissions that are unlikely to be used by other vehicles. In vehicular environments, however, packet collisions still unavoidably occur because *SPS cannot entirely prevent different vehicles from simultaneously selecting the same resources* [14]. This key observation suggests, as we show in this paper, that SPS can be manipulated to launch a DoS attack whose effects are hard to detect among the natural packet loss of the environment.

Contributions: We illuminate the severity of the above vulnerabilities in this paper by crafting novel DoS attacks to exploit them, then we take steps towards improving the security of 5G C-V2X by designing and experimentally validating techniques to detect and mitigate each attack. Our contributions can be summarized as follows:

- We expose the above vulnerabilities in 5G C-V2X by devising, formalizing, and implementing two stealthy DoS attacks that exploit them: (1) *targeted sidelink jamming*, in which an attacker leverages the PHY-layer attributes of C-V2X to anticipate and block up to 93% of BSMs from a target vehicle by jamming just $\leq 10\%$ of the bandwidth of each, and (2) *sidelink resource exhaustion*, in which an attacker abuses SPS operations at the MAC layer to stealthily reduce overall packet delivery ratio (PDR) for vehicles within 1 km range to as little as 22%.
- We showcase the real-world practicality of our DoS attacks through indoor and outdoor experiments using a new portable hardware testbed for 5G C-V2X with software-defined radios and commercial C-V2X devices. We further demonstrate the tangible impacts of our attacks on roadway vehicle dynamics (e.g., severely increased travel times) using traffic simulators.
- We develop lightweight techniques for detecting each attack, using unsupervised machine learning for targeted sidelink jamming and linear regression for sidelink resource exhaustion, and we further propose mitigations for both attacks, for one by leveraging 3GPP Rel-16+ support for aperiodic one-shot transmissions in 5G C-V2X.
- We experimentally validate the effectiveness of our detection and mitigation techniques for both attacks through extensive study in state-of-the-art simulators of 5G C-V2X channels [15] and roadways [16].

The remainder of this paper is organized as follows. In Section II, we provide technical background on the PHY and MAC layers of 5G C-V2X. We state our threat model in Section III and focus each of Sections IV and V on one of our proposed attacks, its detection techniques, and our proposed approaches to mitigate it. We provide our experimental results in Section VI, review related work in Section VII, and conclude with final remarks and future directions in Section VIII.

II. Background

We begin by describing the PHY and MAC layers of the 5G C-V2X protocols: NR-V2X, based on 3GPP Rel-16/17 [17], [18], and LTE-V2X, based on 3GPP Rel-14/15 [19], [20].

A. Sidelink V2V Communication

As a critical safety service, V2V must be supported everywhere, including regions where cellular infrastructure support is limited or nonexistent. Therefore, C-V2X uses *sidelink* communication—in NR Mode 2 or LTE Mode 4—to support direct communication between vehicles. In the absence of base stations (gNBs), 5G C-V2X radios are synchronized using global positioning system (GPS). As per 3GPP Rel-17 [21], LTE and NR sidelink interfaces will both be supported by 5G C-V2X systems for the foreseeable future.

B. C-V2X Sidelink PHY Layers

NR-V2X and LTE-V2X both use 10 ms sidelink frames in the time domain divided into 1 ms subframes (see Fig. 1(a)) [14]. Each subframe is defined as one *Transmission Time Interval (TTI)*, and C-V2X radios record and index PHY/MAC sensing data by TTI. In the frequency domain, C-V2X channels can vary between 10 and 400 MHz wide [14]. However, only 10–20 MHz bandwidths are currently practical (e.g., the U.S. has only allocated 35 MHz for C-V2X [2]). Without loss of generality, we assume a 10 MHz channel in this paper as this is the most common bandwidth and it is supported in both NR- and LTE-V2X. Every C-V2X channel is divided into a number of equal-width logical *subchannels*. The number of subchannels can vary between C-V2X systems; in this paper, we assume the configuration shown in Fig. 1 that is used by our state-of-the-art C-V2X devices (see Section VI): a 10 MHz channel with five 2 MHz subchannels. For convenience, the important notations and abbreviations we use in the remainder of the paper are summarized in Table I.

Each 2 MHz subchannel consists of 10 resource blocks (RBs). As shown in Fig. 1(a), the first two RBs of each subchannel are used for the physical sidelink control channel (PSCCH) and the remainder for the physical sidelink shared channel (PSSCH). Depending on payload size, a transmission may require one or more subchannels. In its lowest-index subchannel (e.g., subchannel 1 for a message spanning subchannels 1–3), a Sidelink Control Information (SCI) message is transmitted over the PSCCH, while the remainder of the allocated subchannel(s) carry a Transport Block (TB) containing the data payload over the PSSCH. Critically, each SCI carries a scrambling code required to demodulate the associated TB and recover the payload [12], [13]; therefore, *failing to recover the SCI results in losing the entire transmission*. We exploit this with the DoS attack presented in Section IV.

C. Semi-persistent Scheduling Algorithm

In sidelink, vehicles use SPS at the MAC layer to autonomously select the TTIs and subchannel(s) they will use to transmit periodic messages (e.g., BSMs) [12], [13]. SPS is mostly identical in LTE- and NR-V2X [14], so we discuss SPS in general with differences noted when they are relevant. SPS is a sensing-based protocol designed to allow short-term prediction of future channel usage based on historical usage data collected during a brief sensing (listening) period. Importantly, C-V2X does not employ any power or multiple-access control mechanisms, so SPS is the *only* means available for vehicles to avert interfering with each other's transmissions. *Semi-persistent* scheduling requires vehicles to periodically *reselect* resources; for BSMs with the standard 10 Hz periodicity, a reselection counter $c \in \{5, \ldots, 15\}$ is randomly set after each reselection. The next reselection is then triggered after c transmissions (i.e., after $c \times 100$ ms).

For a period of time before each resource reselection, SPS requires vehicles to monitor and record channel resource usage. During this *listening* (or *sensing*) period, for TTIs other than those in which the vehicle itself is transmitting, the vehicle records (1) whether it successfully decoded any SCI message, (2) the RBs used by that SCI message and its associated TB (if both are successfully decoded), and (3) the reference signal received power (RSRP) for each RB used by the SCI or TB data [12], [22]. When resource reselection is triggered, a set (denoted by S_A) of candidate single-slot resources (CSSRs) is created to choose new resources from. Each CSSR consists of a number of contiguous subchannels[1] within a single TTI that falls within the *selection window*, defined as a range of future TTIs delimited by TTI offsets T_1 and T_2. In NR-V2X, T_1 and T_2 depend on message priority (specified by the upper layers). For BSMs, the maximum acceptable packet delay is 100 ms, so we assume $T_1 + 1 \leq T_2 \leq 100$. In LTE-V2X, we further have $T_1 \leq 4$. S_A initially contains all possible CSSRs in the selection window, then in a series of steps S_A is reduced to the CSSRs least likely to be used by other vehicles, as follows.

[1]BSMs each require 2 subchannels, based on our observations of C-V2X equipment we use in Section VI.

TABLE I: Important Notations

Abbreviation	Definition
CSSR	Candidate single-slot resource
gNB	Next Generation Node B
GPSDO	GPS-disciplined oscillator
OBU	On-board unit
PDR	Packet delivery ratio
PSCCH	Physical sidelink control channel
PSSCH	Physical sidelink shared (data) channel
RB	Resource block
RRI	Resource reservation interval
RR	Resource reservation
RSRP	Reference signal received power
SCI	Sidelink control information
SPS	Semi-persistent scheduling
TB	Transport block
TTI	Transmission time interval

For each CSSR $R \in S_A$, the last T_0 TTIs of sensing data (where $T_0 = 1000$ in LTE-V2X or $T_0 = 1100$ in NR-V2X) are used to determine whether R should be excluded [14], [22]. The sensed resource R'_k, $k \in \{1, \ldots, \frac{T_0}{100}\}$, which corresponds to the subchannels of R in the $100k$th TTI in the past, is checked. R is excluded if all of the following criteria are true for any R'_k:

1) At least one valid SCI message was decoded from PSCCH in a subchannel within R'_k.
2) For at least one decoded SCI message, a valid TB was received in PSSCH using the resources within R'_k indicated by that SCI message.
3) The average RSRP in the RBs used by a decoded SCI and/or associated TB exceed a threshold th (in dB) that is set based on message priority[2].

Finally, if S_A has been reduced to less than X% of its original size, where $X = 20$ (LTE-V2X) or is a function of message priority (NR-V2X), this process is repeated with th increased by 3 dB. NR-V2X parameterized this requirement and dropped the averaging of RSRP over the sensing window in order to support aperiodic traffic, used for ad hoc V2V and vehicle-to-infrastructure (V2I) services, as well as periodic traffic of more diverse priorities [15]. However, we omit this consideration as we deal only with BSMs, which have equal priority [23].

Upon completion of these steps, the vehicle randomly chooses new resources for its BSMs from the reduced S_A. Critically, this process does not *entirely* prevent vehicles from choosing conflicting resources, it merely attempts to lessen that potential. Multiple vehicles performing simultaneous reselection cannot avoid choosing conflicting resources, vehicles entering communication range after sensing is complete cannot be accounted for, and other peculiarities of V2V may also limit SPS effectiveness. 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. Additionally, as of

[2]This process differs slightly in LTE-V2X [14], but the details of the differences are not relevant to our work.

Releases 16/17, SPS is also used in NR-V2X for scheduling ad hoc aperiodic transmissions. For aperiodic transmissions, the sensing window is reduced to 100 ms; otherwise, SPS is then executed the same way as described above for NR-V2X. We leverage this support for aperiodic transmission in Section IV-C to mitigate one of our attacks.

III. Threat Model for C-V2X DoS Attacks

We consider a single intelligent attacker, Eve, who cares about being efficient and remaining undetected while pursuing different disruptive goals in each attack. We assume she is capable of mimicking an ordinary V2V-equipped vehicle: she may be mobile or stationary and can communicate using sidelink signals on V2V frequencies (e.g., 5.9 GHz band) using portable equipment available for \$3,000 to \$5,000, such as software-defined radios or the commercial off-the-shelf on-board units we use for our indoor and outdoor experiments in Section VI. As we assume Eve wishes to remain stealthy by avoiding detection (and its consequences), we require her to comply with all 5G C-V2X specifications to appear outwardly legitimate:

- Eve must use the standard C-V2X power level of 23 dBm [24] and synchronize to GPS.
- Eve must comply with SPS requirements to regularly reselect resources [12].

When we target a single victim vehicle as in Section IV, we refer to the victim as Alice. We assume Alice implements the most recent V2V security standards (e.g., IEEE 1609.2 [25]) to protect her BSMs and follows all C-V2X requirements (transmits at 23 dBm, sends a BSM every 100 ms, etc.) for normal operation. When we target multiple vehicles, as in Section V, we assume all vehicles in the environment also meet these requirements.

IV. Targeted Sidelink Jamming

In other V2V protocols (e.g., Dedicated Short-Range Communication (DSRC)), it is nearly impossible to predict exactly when a vehicle will transmit a BSM because of medium contention (i.e., random backoff times). However, there is no such contention in 5G C-V2X; SPS is used instead. An attacker need only observe the TTI in which a vehicle transmits *one* BSM to precisely calculate the TTIs of its next several BSMs. Further, the subchannel(s) used to carry the first BSM are also used to carry the subsequent BSMs; therefore, the SCI associated with each BSM will be transmitted in the same two RBs of each corresponding TTI. Through *targeted sidelink jamming*, we show this can be exploited by an attacker to efficiently block each 20-RB BSM from a targeted vehicle by jamming just the associated 2-RB SCI messages. In this attack, we assume Eve can identify the target (Alice) who Eve wants to put at increased risk of collision by preventing Alice's BSMs from being received by other vehicles.

A. Attack Procedure

In order to jam only Alice's BSMs, Eve must be able to differentiate them from others. Direct identification is difficult due to BSM pseudonymization under IEEE 1609.2 [25]; however, commercial standards require BSMs to contain identifying information (e.g., color, make, model) about the sending vehicle [23], and this information is not encrypted. Further, techniques like angle-of-arrival estimation may be employed to isolate Alice's BSMs. Eve executes the following steps to launch the attack against Alice (steps refer to Algorithm 1):

1) *Listen*: In each TTI, Eve receives all BSMs and checks for one from Alice (Step 8).
2) *Record*: Once a BSM from Alice (β_A) is received and identified, Eve marks the TTI (t_A) and SCI subchannel (γ_A) that Alice is using for β_A (Steps 10 – 11).
3) *Anticipate*: Eve calculates the TTIs for Alice's next $c \in \{5, \ldots, 15\}$ BSMs (Step 10).
4) *Jam*: In the calculated TTIs, Eve transmits a narrowband signal (which may or may not be meaningful) to collide only with Alice's SCI field and render the associated TB unrecoverable (as per Section II-C) (Step 5).
5) *Monitor*: Between jamming TTIs, Eve listens for possible BSMs from Alice in unanticipated TTIs (Step 8).
6) *Update*: If monitoring uncovers that Alice has reselected resources (i.e., changed TTI and subchannels), Eve updates her record (Steps 10-11) and then continues jamming Alice's BSMs as before (Step 5).

β_A is an arbitrary BSM from Alice, t_A is the TTI of β_A, and γ_A is the subchannel carrying the SCI message for β_A. Algorithm 1 illustrates the simplicity and efficiency of this procedure. Eve needs to do little more than what an ordinary C-V2X receiver would do; in fact, line 5 in Algorithm 1 is the only abnormal action by Eve that requires meaningful effort.

The critical steps of the attack, *anticipate* and *jam*, are both facilitated by the PHY-layer design of C-V2X. As BSMs are sent precisely every 100 ms, Eve knows that future BSMs from Alice will arrive at 100-TTI intervals after Alice's first BSM, until Alice performs SPS resource reselection. Similarly, the subchannel(s) that Alice uses will remain the same until reselection. Thus, Eve can anticipate and jam Alice's BSMs with extremely high accuracy. If Alice reselects resources and her next BSM arrives earlier than expected, the *update* step ensures Eve will correct herself, failing to jam at most one BSM (with negligible impact on attack effectiveness—

Algorithm 1 Targeted Sidelink Jamming Attack Procedure

```
1:  t_A ← −1
2:  γ_A ← −1
3:  for tti ∈ {0, 1, ..., ∞} do
4:      if tti = t_A then
5:          Transmit jamming SCI in γ_A
6:          t_A = t_A + 100
7:      end if
8:      Receive all BSMs in tti
9:      if β_A ∈ tti then
10:         t_A ← tti + 100
11:         γ_A ← subchannel (β_A)
12:     end if
13: end for
```

see Section VI). In the *jam* step, Eve renders each 20-RB (4 MHz) BSM unrecoverable by jamming only its associated 2-RB (400 kHz) SCI field for a low duty cycle of at most 10%.

B. Attack Detection

In the literature, detecting BSM DoS attacks is generally based on identifying low overall PDR in the C-V2X channel [26]–[29]. However, we show here that such techniques cannot reliably detect our attack. Consider a system monitor of this type that attempts to detect a DoS attack by monitoring the overall PDR in the C-V2X channel for a short time (e.g., 1–2 seconds). We assume the monitor can accurately estimate the number of vehicles (denoted by $\hat{v}$) in its area (e.g., from historical data that correlates vehicle density with time of day [30]) and therefrom estimate the number of BSMs (denoted by $\hat{b}$) that it expects to receive in a given time period using

$$\hat{b} = \hat{v}\lambda t \tag{1}$$

where the length of an observation window in seconds is denoted by t and the rate at which every vehicle transmits BSMs is a constant λ that is known a priori. Then, (1) can be used to devise a PDR threshold Ψ_{th} below which any measured PDR value should raise an attack alarm. The monitor will estimate its observed PDR $\widehat{\Psi}$ over an interval of t seconds as a function of the estimated number of expected BSMs $\hat{b}$ and the number of BSMs that the monitor actually decoded (denoted by b) over the same period. Then, the observed PDR is given by:

$$\widehat{\Psi}(b) = \frac{b}{\hat{b}} = \frac{b}{\hat{v}\lambda t}. \tag{2}$$

Now, this type of monitor can detect the attack (i.e. raise an alarm) if:

$$\widehat{\Psi}(b) < \Psi_{th}(\hat{v}). \tag{3}$$

We note that this formulation is a generic model encompassing a wide range of approaches used in the literature to derive $\Psi_{th}(\hat{v})$. In C-V2X, a number of BSMs are unavoidably lost in packet collisions due to SPS (denoted by α_{SPS}) and, in our model, more packets are lost due to jamming (denoted by α_{eve}). Further, some number of BSMs (denoted by α) will be lost due to communication errors. Thus, we can express b in terms of the actual number of BSMs that were transmitted—which is calculated from the actual (unknown) number of vehicles present (denoted by v)—and these packet losses as:

$$b = v\lambda t - \alpha - \alpha_{SPS} - \alpha_{eve}. \tag{4}$$

Combining (2) through (4) yields a final detection expression

$$\frac{v\lambda t - \alpha - \alpha_{SPS} - \alpha_{eve}}{\hat{v}\lambda t} < \Psi_{th}(\hat{v}) \tag{5}$$

which illustrates why PDR over time is not a reliable metric for DoS detection in C-V2X. Due to the randomness of α and α_{SPS} as well as likely estimation error in $\hat{v}$, any reasonable setting for $\Psi_{th}(\hat{v})$ must have a *significant* error tolerance in order to avoid constantly raising false alarms. This creates an opening for an intelligent attacker, who can carefully constrain α_{eve} by limiting the duration of her attack, to keep the impact on overall system PDR within the error margin of (5). Therefore, the type of PDR-based detection mechanism described by (3)–(5) is unable to reliably detect the targeted sidelink jamming attack.

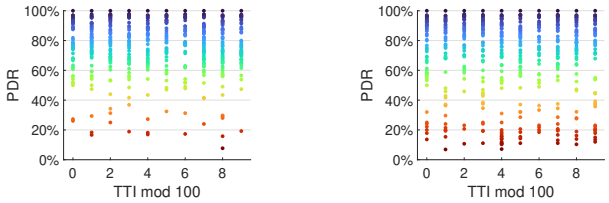

(a) Normal operation.

(b) During targeted sidelink jamming.

Fig. 2: Intuition behind detecting targeted sidelink jamming. TTIs in which jamming occurs (in (b)) have lower relative PDR measurements, facilitating detection with unsupervised clustering.

Our attack has little impact on PDR across many TTIs, but it substantially degrades PDR within *specific* TTIs. While α_{SPS} occurs randomly across all TTIs, α_{eve} occurs only in the TTIs that Alice uses. Fig. 2, which plots per-TTI PDR measurements by a stationary roadside monitor during the simulations described in Section VI-A, depicts the consequences. When there is no attack (Fig. 2(a)), per-TTI PDR is consistently higher than during targeted sidelink jamming (Fig. 2(b)). In particular, Fig. 2(b) suggests *clustering* of measurements with high PDRs, with more scattering (i.e., more outliers) among measurements with low PDR belonging to attacked TTIs. Hence, we propose applying unsupervised learning—specifically, clustering—to detect targeted sidelink jamming when the number of outliers exceeds a tunable threshold.

We use DBSCAN, a time-tested clustering algorithm [31], due to two specific advantages. First, unlike alternatives like k-means clustering, DBSCAN does not require specifying the number of clusters to form, which is important because we cannot reliably predict the PDR for unjammed TTIs (see (5) above). However, we want to ensure that when TTIs at 100 ms intervals are being jammed, the low PDR measurements in those TTIs should not form a cluster. Since Alice reselects her resources at most every 15 transmissions, we set the *minPts* hyperparameter (see below) to 30, thus ensuring that ≤ 15 consecutive TTIs with low PDR due to jamming cannot form a cluster and will instead be (correctly) considered outliers and possibly raise an attack alarm. The attack may be detected as soon as a cluster is formed within any TTI `mod` 100 and outliers become identifiable (see Fig. 2(b)); therefore, this setting of *minPts* ensures the attack is detectable within 30×100 ms = 3 seconds).

The second advantage of DBSCAN is it intrinsically identifies outliers. A cluster is formed if at least *minPts* measurements occur such that each measurement, when rendered in multi-dimensional space, is within distance ϵ of at least one other measurement (we experimented with varying ϵ to determine its best setting in Section VI). Any measurements that do not fall within a cluster are considered outliers—this occurs as part of the core algorithm, eliminating extra pro-

cessing to identify outliers (e.g., local outlier factor) required for alternative algorithms like k-means. DBSCAN thus has an advantage in computational complexity over many alternatives despite its worst-case complexity of $O(n^2)$ [31], where n is the number of measurements. Further, in our specific adaptation, the complexity actually approximates $O(n)$. Since we constrain our system to form clusters and detect outliers only *vertically* (see Fig. 2), for a given measurement the algorithm only needs to check the two nearest points (one above, one below) to determine whether it falls in a cluster, resulting in roughly $2n$ operations overall (granting some leeway for more operations on occasion, e.g., when $minPts$ is reached and a cluster initially forms). As $O(n)$ is efficient enough to run on resource-constrained hardware, we consider DBSCAN to be a more practical choice for our use case.

C. Attack Mitigation

Targeted sidelink jamming is very difficult to mitigate because it exploits fundamental properties of C-V2X: its rigid, slot-based PHY layer and the precise periodicity of BSM. A naive mitigation would simply shorten the time between resource reselections (i.e., reduce c in SPS—see Section II-C) to reduce the number of BSMs that Eve could anticipate and jam based on one BSM from Alice. However, this has been shown to significantly degrade C-V2X performance [32]. A slightly better mitigation would be inducing variations in the periodicity of BSMs, as we initially suggested in [11]. This totally mitigates the attack, as Eve has no way of anticipating future BSMs if she does not know the TTI interval between them, but we subsequently observed that in some scenarios, this approach can negatively impact C-V2X throughput by making SPS less effective at predicting channel use. Therefore, we propose a new mitigation that is able to neutralize the attack *without* degrading SPS performance.

To this end, we propose leveraging the new (as of 3GPP Rel. 17) NR-V2X support for aperiodic transmissions. Both LTE- and NR-V2X support BSM retransmission for more reliable communication. In LTE-V2X, SPS selects fixed resources (in different TTIs) for both an initial and automatic repeat transmission. Targeted sidelink jamming, though, is as effective against LTE-V2X retransmissions as against the original BSMs, because the retransmission resources can be similarly identified and jammed. However, NR-V2X supports ad hoc, *aperiodic* retransmissions that are not quite so easily attacked. By default, NR-V2X devices can inform each other (via a dedicated feedback channel) when a message is incompletely received, allowing the corrupted message to be sent again as a one-time retransmission. This works very well for dealing with packet loss due to a noisy channel or SPS conflicts, but the default configuration still does not mitigate our attack. In order for a vehicle to know a message has been incompletely received, it must receive and successfully decode a valid SCI message, then note that the associated TB cannot be successfully decoded, and finally request retransmission. In our attack, however, the SCI message in each of Alice's BSMs is irrecoverably corrupted by Eve's jamming. Therefore, no SCI can be decoded, so other vehicles will not even know they are missing Alice's BSMs, and the default NR-V2X retransmission behavior will likely never be triggered.

While the default configuration for retransmissions in NR-V2X does not mitigate our attack, we can leverage its support for aperiodic retransmissions to devise an effective mitigation. First, we propose making NR-V2X retransmission mandatory for *all* BSMs; further, we propose treating *every* retransmission as a one-shot aperiodic transmission. Each time a vehicle transmits a BSM, it will use the SPS procedure for aperiodic transmissions (see Section II-C) to select one-time resources and use them to re-transmit that BSM shortly after the original transmission. Unlike LTE-V2X retransmissions, this makes the TTI and subchannel(s) used by each retransmission completely unpredictable to an attacker, and unlike current NR-V2X, this avoids the need for vehicles to receive an SCI message in order to trigger retransmission. Therefore, even if the initial BSM transmissions are jammed by Eve, the retransmissions will still be received by other vehicles, functionally neutralizing targeted sidelink jamming.

Despite the benefit of this approach from a security perspective, we must also consider its cost, as one might reasonably ask whether the increased number of aperiodic transmissions might have a detrimental effect on channel utilization or SPS. However, we note that vehicles can use the *resource reservation interval* (RRI) field of SCI messages to learn whether a received message was periodic or aperiodic. Aperiodic traffic can thus be excluded from consideration during SPS resource selections (i.e., the corresponding resources in S_A will be treated as if they were unused). There is still the chance that aperiodic transmissions may sometimes use conflicting resources, but this will be unlikely to occur very often, especially when the channel is not near maximum capacity, and losing a single BSM to such an occurrence is inconsequential since C-V2X already accepts losing a small number of BSMs (see (5) above). As we show in Section VI-A, this mitigation is completely effective and has negligible collateral impact on system performance. Moreover, we note that retransmitting every message is already supported by the LTE-V2X and NR-V2X standards, so any potential negative impact as a result (if any) is de facto acceptable in current V2V systems. Finally, we note that with NR-V2X now over two years old [33] and vanishingly few vehicles already equipped with V2V, the majority of V2V-equipped vehicles in the near future will support NR-V2X (in NR/LTE-V2X hybrid devices [21]) rather than only LTE-V2X. Therefore, we expect this mitigation will be effective for the overwhelming majority of vehicles.

V. Attack 2: Sidelink Resource Exhaustion

In both 5G C-V2X protocols, SPS is used to predict future channel usage based on observations made during a short listening period. SPS works very well for its intended purpose—minimizing packet collisions due to vehicles choosing the same resources—*if and only if* vehicles largely transmit only periodic BSMs at consistent intervals[3]. SPS is ill-equipped to handle abnormal vehicle behaviors; e.g., a vehicle that transmits periodically, but at varying intervals. If an attacker

[3]With the exception of aperiodic messages in NR-V2X.

exploits this by engaging in such behavior while complying with all other 5G C-V2X specifications, then SPS will not perform its purpose adequately and, as we will show, the C-V2X system can be crippled. In this *sidelink resource exhaustion* attack, we assume Eve wants to cause a DoS effect in the C-V2X channel without *directly* jamming messages from other vehicles. We show how she can act abnormally, while still complying with all C-V2X specifications, to massively degrade C-V2X throughput across a wide geographic area.

A. SPS Vulnerability

When a vehicle performs SPS resource reselection, it aims to identify available 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 all vehicles have the same BSM periodicity (e.g., 100 ms), so a vehicle that sends a BSM in one specific TTI can be expected to transmit subsequent BSMs at 100-TTI intervals for the immediate future. However, since SPS implicitly assumes all transmissions will follow this pattern—with the exception of aperiodic transmissions—it is vulnerable to an attacker who does not meet this assumption. For example, during the listening period, an attacker may transmit several times in one subchannel at 10-TTI intervals, then switch to another subchannel, and still a third, all within the 100-TTI interval when a normal vehicle only transmits one BSM. A listening vehicle running SPS cannot tell that one attacker has made several bogus 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 NR- and LTE-V2X, the SPS listening period is at least 1000 ms, ten times longer than the selection window (see Section II-C). Therefore, *the inclusion of any particular radio resource in S_A is determined based not on one, but on several (at least* 10*) different radio resources* observed during the listening period, any subset of which could be manipulated by an attacker. SPS's lack of perceptual granularity (i.e., the dependence of each CSSR on more than one resource in the listening period) constitutes an exploitable vulnerability in the MAC layer of both 5G C-V2X protocols.

B. Attack Designs

Sidelink resource exhaustion can be executed in different ways. We describe two variants:

1) Variant I: Randomization of legitimate transmissions: In the first variant, Eve generally follows normal C-V2X procedures (i.e., transmitting periodic messages and regularly changing her resources), but we let her control SPS parameters that are fixed or deterministic for normal vehicles. This is formally described in Algorithm 2, where δ_{tti} indicates the periodicity of Eve's transmissions, Γ represents the set of subchannels in the system, γ indicates the size in subchannels of Eve's transmissions, β represents the message contents, and c is the SPS resource reselection counter (see Section II-C). Note that $\mathcal{U}\{a, b\}$ indicates the discrete uniform distribution between a and b. The global channel configuration dictates Γ (e.g., $\Gamma = 5$ in a typical 10 MHz C-V2X channel), but we let Eve control δ_{tti}, γ, β, and the interval from which c is selected. These parameters can be set randomly and tuned over time, so Variant I does not require any preparation time before executing the attack. The parameters can also be set based on specific observations of the environment (e.g., roadway vehicle density, noise levels) to be able to adapt to different environments. Eve will then transmit c messages periodically and change her resources like a normal vehicle (although her messages will contain bogus data rather than an authentic, i.e., identifiable, BSM). This variant of the attack is simple, but extremely effective. Eve can transmit only $2-5$ times as often as a regular vehicle (determined by her selection of δ_{tti}), and in doing so she adds perturbations into SPS that, as we show in Section VI-B, can ultimately reduce PDR significantly for all vehicles within several hundred meters.

Eve sets δ_{tti} to whatever periodicity she desires and she may choose to vary this parameter over time. Like a normal vehicle, she will periodically transmit messages–in her case, filled with bogus or randomly generated data rather than an authentic (i.e., potentially identifiable) BSM–and change her resources after c transmissions. Unlike other vehicles, however, Eve can choose her own range of values for c. Furthermore, we allow Eve to deliberately select her new resources based on whatever factors she chooses rather than based on the typical SPS listening process. For example, she may set δ_{tti} for her transmissions and settle on a discrete interval $\{c_{min}, c_{max}\}$ from which to randomly select values of c, then execute the attack as outlined in Algorithm 2 by randomly setting other parameters as needed. This variant of the attack is fairly unsophisticated, but extremely effective. Just by adding random perturbations into SPS data, we show in Section VI-B that this attack can reduce PDR for all vehicles within several hundred meters of Eve to as little as 40%.

2) Variant II: Exploiting SCI for predictive abuse of SPS: Every SCI message in both C-V2X protocols contains a resource reservation (RR) field, a single-bit field that tells receivers whether the resources used by that SCI message and its associated TB will be used for another periodic transmission by the same vehicle [13]. This information is an important decision factor when vehicles select new resources [22]; unfortunately, it also gives critical information to an attacker. If Eve listens to the channel for just 100 ms (one BSM interval) and records the RR value for every decoded SCI, she can predict with high accuracy which resources in the next BSM interval will be used by other vehicles. Then, she can tailor her transmissions to achieve a negative impact on the system *in combination with those of other vehicles*, thus making every other vehicle in the system an unwilling, passive participant in the attack. For example, if Eve listens for 100 ms and observes

when vehicles use a particular subchannel or subchannels, she can transmit strategically to "fill in" those subchannels over the next 100 ms (or longer). Other vehicles will then register her transmissions in combination with other vehicles' as indicating the filled subchannels cannot be used, and they will most likely select resources in other subchannels. Fig. 3 shows a scaled-down example of this. In the depicted situation, three strategic transmissions by Eve reduce the number of candidate resources from 11 to 5 - a reduction of 54% from just a 15% duty cycle for Eve. In Section VI-B, we show that this variant of sidelink resource exhaustion nearly doubles the impact of Variant I at a cost of just 100 ms of initial preparation time, reducing overall packet delivery ratio in the C-V2X channel to as little as 22%.

C. 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 an abnormal decrease in channel PDR; however, inferring the cause is far more difficult. As Eve does not deliberately jam BSMs directly, there is no observable correlation of her transmissions with lost packets. Further, Eve's relatively normal (albeit modified) C-V2X operation will not noticeably stand out from normal traffic. Although δ_{tti} is likely to be less than a normal vehicle's rate of BSM transmission, the anonymous nature of V2V transmissions makes identifying this anomaly extremely difficult (e.g., three transmissions from Eve cannot be distinguished from single transmissions by three vehicles surrounding her). Further, the size γ of Eve's messages is consistent with ordinary traffic (i.e., there is no broadband jamming in any TTI) and her transmit power is the same as for any other vehicle. Therefore, we contend that attempting direct detection of Eve by identifying her transmissions as malicious is unlikely to succeed.

Instead, we propose that a monitor should watch channel resource utilization 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 in the aftermath of the attack; i.e., a large percentage of channel bandwidth will be wasted. This narrowing effect does not occur under normal operating conditions and can therefore be considered as a sort of attack signature for sidelink resource exhaustion. By monitoring channel resource utilization patterns, it is possible to observe when the number of used resources diminishes unexpectedly for a period of time, potentially facilitating detection of the attack. One way to monitor channel resource utilization in this fashion is to approximate trends using least-squares regression analysis. We demonstrate the effectiveness of this approach in Section VI-B2.

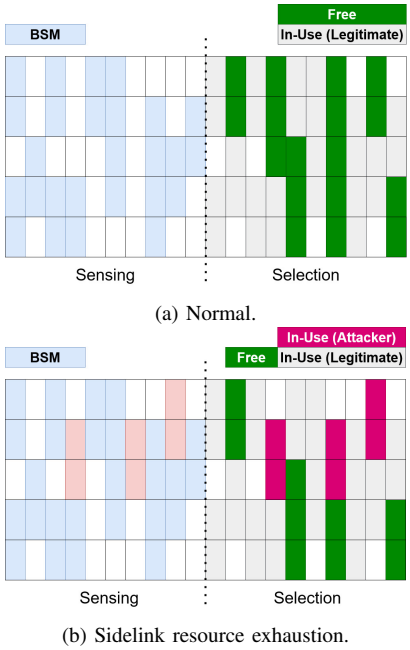

(a) Normal.

(b) Sidelink resource exhaustion.

Fig. 3: A scaled-down example of SPS results with 10 ms sensing and selection windows for benign (a) and sidelink resource exhaustion variant II (b).

Algorithm 2 Sidelink Resource Exhaustion Attack Procedure

Require: $|\Gamma| > 0$
$\delta_{tti} \leftarrow 10$
$\beta \leftarrow$ random data
$\gamma \leftarrow \mathcal{U}\{1, |\Gamma|\}$
$|\beta| \leftarrow \gamma$
$c \leftarrow \mathcal{U}\{c_{min}, c_{max}\}$
for $tti \in 0, 1, \ldots, \infty$ **do**
  if $tti \bmod \delta_{tti} = 0$ **then**
    Transmit β in subchannels $(\gamma, \gamma + |\beta|)$
    $c = c - 1$
  end if
  if $c = 0$ **then**
    $\gamma \leftarrow \mathcal{U}\{1, \Gamma\}$
    $|\beta| \leftarrow \gamma$
    $c \leftarrow \mathcal{U}\{c_{min}, c_{max}\}$
  end if
end for

D. Attack Mitigation

Our sidelink resource exhaustion attack exploits the disparity in size between the listening period and the set of CSSRs S_A assembled during SPS, which we contend constitutes a fundamental vulnerability in the design of the SPS algorithm. To mitigate the attack, this vulnerability must be resolved. As two sides of one approach, either the size of S_A could be expanded or the length of the listening window could be reduced, both of which would remove the size disparity exploited by our attack. The former option is infeasible *prima facie*, because it would require increasing the T_2 parameter in SPS (see Section II-C) to 1000 or 1100 ms in LTE- or NR-V2X, respectively. This would potentially allow for a full second to pass between BSMs during resource reselection, violating the 100 ms latency requirement for safety-critical messages. Therefore, we focus our attention on the latter option, reducing the length of the SPS listening period from

TABLE II: Simulation Parameters for Targeted Sidelink Jamming Evaluation

Roadway	Road length	2 km
	Road width	24 m
	Number of lanes	6 (3 each direction)
	Average number of vehicles present	71 ± 2
	Vehicle speeds (kph)	$\mathcal{N}(100, 10)$
	Vehicle density	35 vehicles/km
Communication	Alice's transmit power	23 dBm
	Eve's transmit power	23 dBm
	Modulation	16-QAM
	SINR threshold to decode SCI	−1000 dBm
	SCS (NR)	15 kHz
	Bandwidth	10 MHz
	Number of subchannels	5×2 MHz
	Channel Model	3GPP 5G Model

1000 or 1100 ms to a shorter period that is closer to the size of S_A. The idea here is to eliminate the influence of earlier, less relevant transmissions (including the attacker's) on SPS resource selection; in fact, Rel-17 NR-V2X does this very thing, but only for aperiodic messages [14]. In Section VI, we show that reducing the listening period to 100 ms effectively mitigates the attack with no significant negative impact on C-V2X system performance.

VI. Experimental Evaluation

To evaluate our work, we first constructed a prototype hardware C-V2X testbed, using USRP software-defined radios and commercial C-V2X evaluation kits, to support over-the-air experiments that demonstrate the real-world practicality of our work. Our hardware experiments further corroborate our simulation results (see below) and support the conclusions we draw from them.

In our testbed, USRP B210s with one or two 5 dBi antennas and TCXO GPSDO modules (for GPS time synchronization) emulate vehicles. Each USRP is connected to one laptop running our significantly extended[4] version of the srsRAN [35] project, and each laptop/USRP pair can emulate one or more vehicles. Beyond indoor experiments (e.g., as configured in Fig. 4(a)), our testbed is mobile and distributable, and it can even be deployed in real vehicles (like in Fig. 6(a)) for dynamic outdoor experiments. Along with USRPs, in some experiments we use commercial Cohda MK6C LTE-V2X on-board units (OBUs) [36]. Each MK6C is an evaluation kit, featuring an NXP iMX8 chipset and two 4 dBi C-V2X antennas, that is designed for use in a real vehicle. Unlike our USRPs, each Cohda OBU can only emulate one vehicle, so we use them only in experiments where a large number of vehicles is not required (e.g., to evaluate targeted sidelink jamming). For GPS synchronization of all devices, we either use real GPS satellite signals (in outdoor experiments) or we synthesize GPS signals using *gps-sdr-sim* [37] and transmit them over-the-air inside our RF-shielded laboratory using the LimeSDR shown in Fig. 4(a).

Beyond hardware experiments, we evaluate our attacks, detection techniques, and mitigations using two state-of-the-art simulators—WiLabV2Xsim [15] and VEINS [16]—to study impacts on the V2V network and roadway traffic dynamics, respectively. WiLabV2Xsim [15] is a widely used V2V simulator that features 3GPP-compliant implementations of both NR-V2X and LTE-V2X. Simulations in WiLabV2Xsim allow us to evaluate our work using NR-V2X (for which no commercial devices yet exist); consider realistic vehicle densities, speeds, etc. that cannot be easily replicated in hardware experiments; and evaluate our work in realistic, measurement-based V2V channels. We note that while we performed all WiLabV2Xsim experiments using both NR- and LTE-V2X, we found the results were nearly identical, so for clarity we provide only NR-V2X results for all experiments. For sidelink resource exhaustion, we use VEINS, another well-established simulator, to study how the effects of the attack produce tangible results in terms of traffic flow and vehicle dynamics. We do not consider VEINS for targeted sidelink jamming as the intent of that attack is not to impact a large number of vehicles.

A. Targeted Sidelink Jamming

We first evaluate targeted sidelink jamming in hardware experiments and then in simulations.

1) Attack: In our hardware testbed, we configure two Cohda OBUs, which each transmit BSMs at the standard 10 Hz rate, to represent two vehicles. We set these "vehicles" (Alice and Bob) 1 m apart while Eve, a two-antenna USRP B210, is placed 2 m from both Alice and Bob. We let Alice and Bob run normally (without an attacker) for 10 minutes to establish a baseline PDR; by comparing Bob's received packet log against Alice's packet transmission log, we find Bob received >99.85% of Alice's BSMs, indicating a very low rate of packet loss with no attacker. Next, we repeat this experiment, but we add the attacker, Eve, who executes targeted sidelink jamming attack against Alice. After 10 minutes, we again compare Alice's and Bob's packet logs; as shown in Fig. 4(b), this reveals the attack was extremely successful. Within two seconds of the attack beginning, Alice's PDR drops below 10% and remains there for the duration of the experiment (Fig. 4(b) only shows the first 10 seconds to emphasize the time frame when the attack effects begin—Alice's PDR remains consistently low thereafter). This clearly illustrates the devastating impact on Alice's ability to inform other vehicles of her presence, which significantly increases her chance of colliding with another vehicle. Further, Fig. 4(b) shows the negligible impact on Bob, demonstrating the minimal collateral damage and, consequently, the stealthiness of the attack. Finally, we note these results prove the vulnerability of commercial C-V2X devices to targeted sidelink jamming. If an attacker obtains the necessary equipment (see Section III), state-of-the-art commercial devices can be easily attacked.

Using WiLabV2XSim, we now evaluate the effectiveness of targeted sidelink jamming in a realistic V2V channel and

[4]After enhancing the bare-bones, receiver-only sidelink implementation from srsRAN, we extended and augmented work by Eckermann and Wietfeld [34] to implement a full C-V2X transceiver and further developed code to implement our attacks.

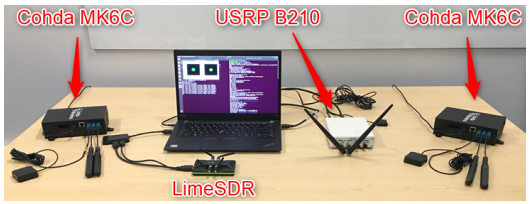

(a) Experimental hardware testbed.

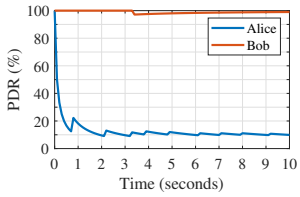

(b) Testbed results.

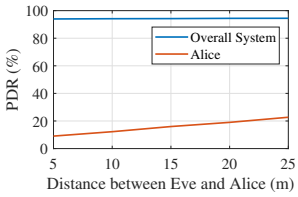

(c) Simulation results.

Fig. 4: Experimental results for targeted sidelink jamming.

observe how the attack impact changes with the physical distance between Eve and Alice. We consider the highway scenario described by Table II and simulate the attack with Eve positioned in the same lane as Alice (e.g., following her in a car) at varying distances from 5 – 25 meters. For each distance in that interval, we run 100 iterations of the scenario for 60 simulated seconds and calculate the average of the results for each distance. From those results, we observe a linear relationship between Eve's distance from Alice and her ability to degrade Alice's PDR (see Fig. 4(c)), with a trade-off of roughly 5% less impact on PDR for every 5 m further Eve gets from Alice. At 5 m, a similar distance as used in our laboratory experiments, Eve can knock out just over 90% of Alice's BSMs, corroborating our hardware results. However, even at 25 m, Eve can jam ~80% of Alice's BSMs, allowing her to devastate Alice's ability to communicate safety-critical information to nearby vehicles while remaining at an inconspicuous distance.

2) Detection with DBSCAN: We begin by determining the best value of the ϵ hyperparameter (see Section IV-B). To do this, we use the parameters in Table II to run simulations with and without an attacker, varying ϵ between 0 and 0.03 (selecting this range based on sorted k-nearest neighbor analysis). For each value of ϵ, we calculate the true positive and false positive rates to create the receiver operating characteristic (ROC) curve shown in Fig. 5(a). Depending on what action is to be taken upon detecting an attack (i.e., on the cost of raising a false alarm), this ROC curve can help a defender select a trade-off between true and false positive rates; however, deciding on a specific action and corresponding trade-off is beyond our scope. Our goal is to show the potential of our approach, and we contend we have done so. Considering especially that in our system model, a detector only has a limited data set (the per-TTI PDR over the tens of seconds that an attacker and victim might be in range of the detector), our 80% true positive rate in exchange for 25% false positive rate is reasonable.

That said, one might reasonably wonder if deep learning could obtain better results. To explore this, we evaluate the potential for using a long short-term memory (LSTM) recurrent neural network to detect the attack using the same

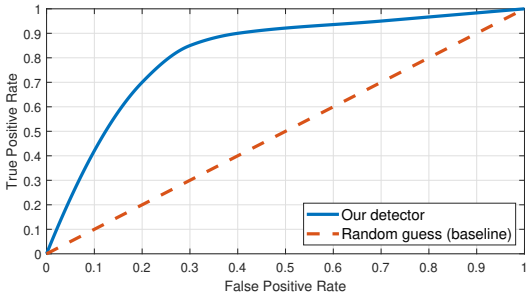

(a) ROC curve for DBSCAN-based detection when varying the ϵ parameter.

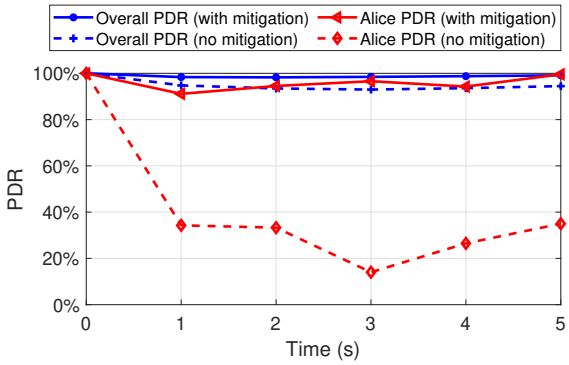

(b) PDR of the NR-V2X system and for Alice with and without our mitigation for targeted sidelink jamming.

Fig. 5: Detection and mitigation results for targeted sidelink jamming.

limited data set as our lightweight DBSCAN approach. Due to the low variance among PDR measurements over time, an LSTM cannot use very many hidden layers without overfitting (our model overfits within one training epoch for ≥ 3 hidden layers). We therefore construct an LSTM with just one hidden layer and a softmax head for classification (i.e., attack detection) and train it on 1000 samples, evenly distributed between the "attack" and "no attack" classes, collected from WiLabV2Xsim simulations. After training, our LSTM obtains a true positive rate of 90% on the test set, but only at the expense of a huge 75.3% false positive rate. In contrast, our DBSCAN approach can obtain the same 90% true positive rate with only a 40% false positive rate (see Fig. 5(a)) using the same simple dataset. Therefore, our approach obtains comparable, and somewhat superior, results than a common deep learning approach, while also being more lightweight and hence more easily deployed on resource-constrained vehicular equipment than deep neural networks.

3) Mitigation: We now evaluate the effectiveness of our mitigation for targeted sidelink jamming (see Section IV-C) by repeating our attack simulations from Section VI-A with the mitigation applied and contrasting the results. We implement the mitigation in WiLabV2Xsim by selecting the option to retransmit every BSM, then modifying the necessary code files so that the resources used for each retransmission are chosen using the SPS procedure for one-shot aperiodic trans-

missions. Our results, shown[5] in Fig. 5(b), are excellent. With the mitigation in place, Eve is no longer able to jam the retransmissions of Alice's BSMs, and Alice's PDR remains over 95% during the attack (vs. ≤ 40% without the mitigation). Also, despite the increased number of aperiodic transmissions, overall PDR remains ≥ 90%, corroborating our expectation (see Section IV-C) of no significant negative impacts on overall C-V2X performance.

B. Sidelink Resource Exhaustion

We previously confirmed that the vulnerability this attack exploits (see Section V-A) is present on commercial C-V2X equipment [11]; now, we seek to establish the feasibility of the attack against C-V2X hardware in an outdoor environment under more realistic conditions. Further, we study the impact on a vehicular network with realistic V2V channels using WiLabV2Xsim. We then use VEINS to study the impact on vehicle dynamics (e.g., increased travel times). Finally, we evaluate the effectiveness of our detection technique and mitigation using WiLabV2Xsim.

1) Attack: For outdoor experiments, we use an experimental setup similar to Fig. 6(a). Using C-V2X transmissions recorded in WiLabV2Xsim for a highway scenario (see Table II) one USRP acts as the victim of the attack while another USRP transmits the traffic for all other vehicles, including the attacker. We evaluate only the first variant of sidelink resource exhaustion in these experiments as the quantity of traffic generated in the second variant exceeds the capacity of our rate-limited, USB-connected USRP B210s.

Because our experiments are conducted in a noisy outdoor environment with imperfect GPS coverage and our USRPs transmit at a lower power level than commercial C-V2X devices (10–12 dBm vs. 23 dBm), overall PDR is low (∼ 42%) even without an attacker over a 30-second period where a cart-mounted USRP moves along a radius 25 m from the victim. However, when the experiment is repeated with an active attacker, the victim vehicle is still severely impacted, with the PDR dropping sharply to just over 5% within 30 seconds. This reaffirms the real-world viability of our attack and shows it can be executed in a more realistic outdoor setting without obtaining expensive commercial OBUs, illustrating the severity of this threat.

We now use WiLabV2Xsim to evaluate both variants of sidelink resource exhaustion in a realistic system with measurement-based V2V channels. We again consider the highway scenario from Table II, except in these experiments we vary the density of vehicles on the road to study how the attack effects scale with the size of the V2V network and level of C-V2X channel busyness. Using government traffic data from July 2017[6] [38], we analyze measurements from a 2 km stretch of the I-490 highway near Rochester, NY and find that the highest average vehicle density (during morning and evening "rush hours") is 89 vehicles/km, for a total of 178

(a) Outdoor experimental prototyping of sidelink resource exhaustion attack.

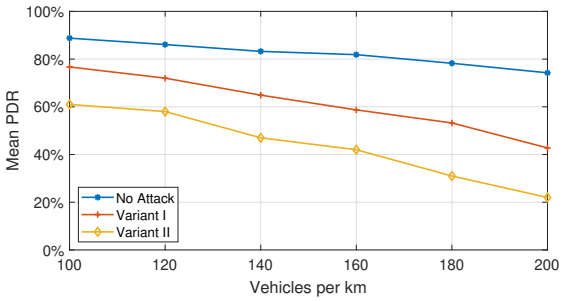

(b) Impact of the sidelink resource exhaustion attack for different vehicle densities in the highway scenario.

Fig. 6: Experimental testbed and simulation results for sidelink resource exhaustion.

vehicles in our 2 km simulated stretch of highway. At the low end, in mid-morning hours, we calculate an average density of 48 vehicles/km (96 total vehicles). Rounding these extremes, we evaluate sidelink resource exhaustion experimentally for vehicle densities between 100–200 vehicles.

We perform Monte Carlo simulations with 100 iterations to evaluate each attack variant. For each iteration, Eve sets her Variant I parameters (see Algorithm 2) to $\delta_{tti} = 20$, $\Gamma = \{1, \ldots, 5\}$, $\gamma \sim \mathcal{U}\{1, 5\}$, $c_{min} = 2$ and $c_{max} = 10$. For Variant II, Eve uses 6 MHz transmissions to attempt to deny the use of subchannels 1–3 to other vehicles (see Section V-B2). Fig. 6(b) shows our results, clearly illustrating the severe, negative impact on overall PDR in the C-V2X channel for both attack variants, at any size V2V network. Even at the lowest vehicle density (100 vehicles), Eve can reduce overall PDR from 90% to 78% (−12 points) with Variant I and to 63% (−27 points) with Variant II, an already severe impact. However, when 200 vehicles are present, Eve can achieve an even harsher effect, reducing PDR from 75% to 43% (−32 points) under Variant I and inflicting a devastating drop from 75% to just 22% (−53 points) under Variant II. Thus, when the system is at or near capacity, the *maximum* PDR in the C-V2X channel under either attack variant is 42%, which is completely unacceptable in a safety-critical application. This further shows our advantage; e.g., at a vehicle density of 150, the attacks in [10] only reduce packet reception in the network to about 93% [10, Fig. 16], whereas sidelink resource exhaustion at the same vehicle density can reduce PDR to just 60% (Variant I) or 42% (Variant II).

Finally, we consider how the effects of our attack may impact the dynamics of vehicles on the road by emulating

[5]For clarity, we show only the results for 15 m distance between Eve and Alice—the results are similar for all distances.

[6]We use 2017 data as it is the most recent available data except for April 2020. The 2020 data is a poor reference point because a Covid-19 lockdown at that time created abnormally low traffic levels throughout the state.

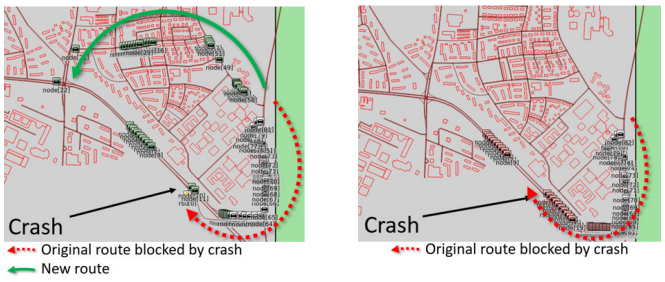

(a) Normal V2V operation.

(b) During sidelink resource exhaustion.

Fig. 7: Simulation results from VEINS illustrating impact of sidelink resource exhaustion.

those effects (as measured in WiLabV2Xsim) using VEINS. As no public data exists on how vehicles or drivers may make decisions based on information received via V2V, we make a few reasonable assumptions: (1) any vehicle that is stopped in traffic unexpectedly (e.g., not at a controlled intersection) for more than 15 seconds will include an alert in its next BSM about an excessively congested roadway, and (2) any vehicle that receives such an alert will choose to re-route around the blocked road, if an alternative route is available. With these assumptions, we simulate an urban scenario where 40–50 vehicles are traveling in one direction on a major highway when a crash occurs, obstructing the road for approximately 3 minutes, and observe what happens when we do and do not introduce the effects of the attack. As Fig. 7(a) shows, under normal circumstances, the majority of vehicles get re-routed after receiving an alert via V2V and continue on their way, whereas during the attack, the majority of vehicles do *not* re-route and are instead subjected to at least a 3-minute delay while the crash is cleared from the roadway. In concrete terms, we calculate an average additional wait time of 2.4 minutes during the attack, demonstrating its tangible impacts on vehicle dynamics—and this is a scenario with just $40-50$ vehicles. On a busy highway with denser traffic, these delays will undoubtedly be even more significant, and the gridlock resulting from longer delays will take longer to clear, compounding the severe effects of our attack.

2) Detection through regression analysis: We use WiLabV2Xsim to evaluate our technique for detecting sidelink resource exhaustion using least-squares regression analysis of channel resource utilization over time. Our results are similar for both attack variants, so we focus just on Variant I for clarity. For continuity, consider the same scenario described above for evaluating sidelink resource exhaustion in a highway environment. We assume a stationary monitor positioned midway along our 2 km stretch of highway who records channel resource utilization over time (e.g., using channel busyness ratio). We run the simulations with and without the attack to compare the results, which are shown in Fig. 8(a). When there is no attack, channel usage remains roughly constant at 70%. During the attack, however, it takes less than ten seconds for channel utilization to fall below 40% as vehicles react to the attack by using less bandwidth, and the downward trend in utilization is noticeable within 2 seconds. Applying

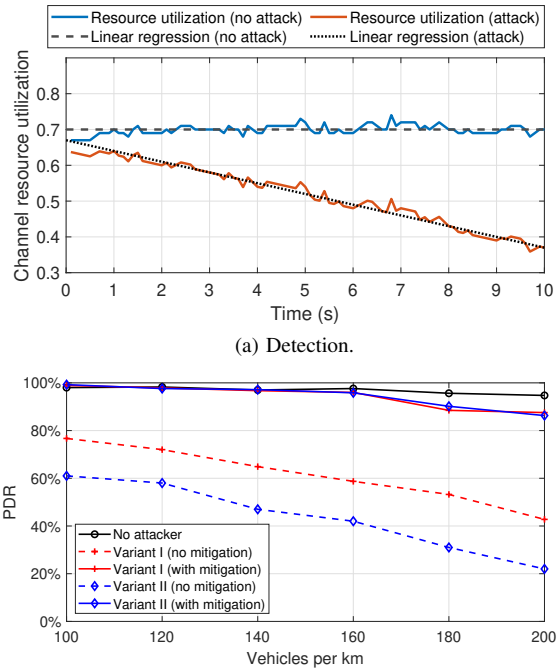

(a) Detection.

(b) Mitigation.

Fig. 8: Detection and mitigation results for sidelink resource exhaustion.

least-squares regression, we find that our monitor observes channel utilization decrease at a rate of $-0.03t$ (where t is the elapsed time, in seconds, since the start of the attack) during the attack, as compared with no measurable change when there is no attack. We contend this will allow the monitor to raise an alarm when either channel utilization trends downwards, or utilization falls below an application-defined threshold, depending on the needs of particular use cases.

3) Mitigation: We evaluate our mitigation for sidelink resource exhaustion (see Section V-D) by again using WiLabV2Xsim experiments. To do this, we repeat the attack experiments described above with the mitigation applied, varying the shortened SPS listening period between $50-700$ ms. The results for different periods vary with no clear trend. Our best results, shown in Fig. 8(b), occur when the listening period is shortened to 100 ms, coincidentally matching the size of the set of CSSRs (S_A from Section II-C). As Fig. 8(b) makes clear, shortening the SPS listening period to 100 ms substantially reduces the effectiveness of both sidelink resource exhaustion variants. While Eve still has a noticeable impact on PDR at the highest vehicle densities even with the mitigation applied, she is not able to reduce PDR below 85% (as opposed to reducing it as low as 20–30% without the mitigation). This clearly illustrates the effectiveness of this mitigation and further demonstrates its negligible collateral effects, supporting the adoption of this approach in future 3GPP standards.

VII. Related Work

DoS attacks at the PHY and MAC layers have been extensively studied against DSRC (e.g., [27], [28], [39]), but

such work is generally inapplicable to C-V2X due to the huge differences between the protocols. To take just one example, Hussein *et al.* [39] designed a DoS attack against DSRC that exploits the back-off timer used for congestion control in that protocol; as C-V2X has no similar mechanism (see Section II), such an attack is not possible against C-V2X. To the best of our knowledge, only two prior works describe DoS attacks at the PHY- or MAC-layer that specifically target C-V2X. First, Trkulja *et al.* [10] presented an attack where one or more attackers collaboratively exploit SPS to anticipate and jam BSMs from other vehicles. The intuition behind this is similar to our targeted sidelink jamming attack. However, our novel attack is easier to execute, in part because we only require one attacker, whereas [10] requires multiple collaborating attackers to have a meaningful effect. Further, our attack targets a *specific* victim while theirs impacts other vehicles arbitrarily. Finally, our system model more completely reflects C-V2X standards than [10] does, and instead of a simplified LTE-V2X simulation environment we have evaluated our attacks in a standard-compliant 5G C-V2X simulator *and* on hardware using state-of-the-art LTE-V2X evaluation kits. Second, Li *et al.* [9] proposed a resource exhaustion attack based on an intuition somewhat similar to our sidelink resource exhaustion attack; however, they targeted infrastructure-based resource allocation in LTE sidelink Mode 3, whereas we target autonomous resource selection by vehicles in sidelink Mode 4. Further, our attack (unlike [9]) cannot be mitigated by network-layer filters, and our attack is much less detectable because our attacker complies with C-V2X standards and appears "normal" while theirs must significantly deviate from typical behaviors.

BSM DoS attack detection is often based on monitoring PDR [27]–[29]. 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 DSRC detection techniques generally inapplicable to C-V2X.

VIII. Conclusion

5G C-V2X promises to deliver the long-promised safety benefits of V2V as well as myriad other improvements to the transportation experience, but the effectiveness of 5G C-V2X 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 5G C-V2X protocols. 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 evaluated mitigations for both attacks, thus providing directions for improving the security of 5G C-V2X in future releases of 3GPP standards.

References

[1] 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

[2] Federal Communications Commission, "Use of the 5.850–5.925 GHz band," 35 FCC Rcd 13440 (16), Nov. 2020.

[3] *LTE-V2X Access layer specification for Intelligent Transport Systems operating in the 5 GHz frequency band*, ETSI European Standard 303 613, Rev. 1.1.1, Jan. 2020.

[4] National Highway Traffic Safety Administration (NHTSA), "Notice of Proposed Rulemaking: Federal Motor Vehicle Safety Standards; V2V Communications," 82 Fed. Reg. 3854, pp. 3854–4019, Jan. 2017.

[5] Y. Liu *et al.*, "Secrecy rate maximization via radio resource allocation in cellular underlaying V2V communications," *IEEE Trans. Veh. Technol.*, vol. 69, no. 7, pp. 7281–7294, Apr. 2020.

[6] 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.

[7] 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.

[8] 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, Apr. 2020.

[9] 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)*, Abu Dhabi, United Arab Emirates, Dec. 2018.

[10] N. Trkulja, D. Starobinski, and R. A. Berry, "Denial-of-service attacks on C-V2X networks," in *Proc. Int. Workshop on Automotive and Autonomous Vehicle Security (AutoSec)*, Virtual, Feb. 2021.

[11] G. Twardokus and H. Rahbari, "Vehicle-to-nothing? Securing C-V2X against protocol-aware DoS attacks," in *Proc. IEEE Int. Conf. on Computer Communications (INFOCOM)*, Virtual, May 2022, pp. 1629–1638.

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

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

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

[15] V. Todisco, S. Bartoletti, C. Campolo, A. Molinaro, A. O. Berthet, and A. Bazzi, "Performance analysis of sidelink 5G-V2X Mode 2 through an open-source simulator," *IEEE Access*, vol. 9, pp. 145 648–145 661, Oct. 2021.

[16] C. Sommer, R. German, and F. Dressler, "Bidirectionally coupled network and road traffic simulation for improved IVC analysis," *IEEE Trans. Mobile Comput.*, vol. 10, no. 1, pp. 3–15, Jan. 2011.

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

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

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

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

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

[22] *Physical layer procedures for data (Release 17)*, 3GPP Technical Specification 38.214, 2022.

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

[24] *User Equipment (UE) radio transmission and reception*, 3GPP Technical Report 36.785, 2016.

[25] *Wireless Access in Vehicular Environments (WAVE)–Security Services for Applications and Management Messages*, IEEE Standard 1609.2-2016, 2016.

[26] 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.

[27] 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.

[28] ——, "AI-Based malicious network traffic detection in VANETs," *IEEE Netw.*, vol. 32, no. 6, pp. 15–21, Nov. 2018.
[29] 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, pp. 157–162.
[30] 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, Jul. 2021, pp. 129—-134.
[31] E. Schubert, J. Sander, M. Ester, H. P. Kriegel, and X. Xu, "DBSCAN revisited, revisited: Why and how you should (still) use DBSCAN," *ACM Trans. Database Syst.*, vol. 42, no. 3, pp. 1—-21, Jul. 2017.
[32] 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.
[33] *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.
[34] 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.
[35] 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.
[36] Cohda Wireless, "MK6c EVK - Cohda Wireless," 2020, Accessed: Oct. 22, 2022. [Online]. Available: https://www.cohdawireless.com/solutions/hardware/mk6c-evk/
[37] T. Ebinuma, "GPS-SDR-SIM," 2018, Accessed: Apr. 20, 2020. [Online]. Available: https://github.com/osqzss/gps-sdr-sim
[38] "Traffic Data Viewer: Report for station 542 (43.14168,-77.5904) for week of July 10, 2017," New York State Department of Transportation, 2022, Accessed: Apr. 17, 2022. [Online]. Available: https://www.dot.ny.gov/tdv
[39] S. Hussein, M. S. Mohamed, and A. Krings, "A new hybrid jammer and its impact on DSRC safety application reliability," in *Proc. IEEE Annu. Inf. Technol. Electron. Mobile Commun. Conf. (IEMCON)*, Vancouver, BC, Canada, Oct. 2016.

Geoff Twardokus (Graduate Student Member, IEEE) received the M.S. and B.S. degrees in computing security from Rochester Institute of Technology (RIT), NY, USA in 2021. He is currently pursuing the Ph.D. degree in Electrical and Computer Engineering at RIT. His research interests include security in emerging and next-generation vehicle-to-everything (V2X) technologies, physical-layer security in wireless networks, and post-quantum security for connected vehicles.

Hanif Rahbari (Member, IEEE) received the PhD degree in electrical and computer engineering from the University of Arizona in 2016. He is currently an assistant professor with the Golisano College of Computing and Information Sciences and a member of ESL Global Cybersecurity Institute and Electrical and Computer Engineering PhD program, Rochester Institute of Technology (RIT). He joined RIT in Spring 2018 after a brief experience as a post-doctoral associate with Virginia Tech. His broad research area is wireless ecosystem security and wireless communications, with emphasis on jamming and privacy (transmission attributes obfuscation) at the physical layer, connected vehicles security, applied cryptography in wireless systems, and spectrum sharing security. He received the NSF CAREER Award in 2023 and is a co-inventor on three US patents.