7-30-2021

Intelligent Lower-Layer Denial-of-Service Attacks Against Cellular Vehicle-to-Everything

Geoff Twardokus
gdt5762@rit.edu

RIT | Golisano College of Computing and Information Sciences
Department of Computing Security

Rochester Institute of Technology

Golisano College of
Computing and Information Sciences

Master of Science in

Computing Security

Thesis Approval Form

Student Name: Geoff Twardokus

Thesis Title: Intelligent Lower-Layer Denial-of-Service Attacks
Against Cellular Vehicle-to-Everything

Thesis Committee

Name Signatures Date

Hanif Rahbari
Committee Chair

Sumita Mishra
Committee Member

Amlan Ganguly
Committee Member

Intelligent Lower-Layer Denial-of-Service Attacks Against Cellular Vehicle-to-Everything

Geoff Twardokus

Committee Members:

Hanif Rahbari

Sumita Mishra

Amlan Ganguly

A thesis submitted in partial fulfillment of the requirements for the degree of Master of Science in Computing Security

Rochester Institute of Technology
Golisano College of Computing & Information Sciences
Department of Computing Security

July 30, 2021

Abstract

Vehicle-to-everything (V2X) communication promises a wide range of benefits for society. Within future V2X-enabled intelligent transportation systems, vehicle-to-vehicle (V2V) communication will allow vehicles to directly exchange messages, improving their situational awareness and allowing drivers or (semi-)autonomous vehicles to avoid collisions, particularly in non-line-of-sight scenarios. Thus, V2V has the potential to reduce annual vehicular crashes and fatalities by hundreds of thousands. Cellular Vehicle-to-Everything (C-V2X) is rapidly supplanting older V2V protocols and will play a critical role in achieving these outcomes. As extremely low latency is required to facilitate split-second collision avoidance maneuvers, ensuring the availability of C-V2X is imperative for safe and secure intelligent transportation systems. However, little work has analyzed the physical- (PHY) and MAC-layer resilience of C-V2X against intelligent, protocol-aware denial-of-service (DoS) attacks by stealthy adversaries. In this thesis, we expose fundamental security vulnerabilities in the PHY- and MAC-layer designs of C-V2X and demonstrate how they can be exploited to devastating effect by devising two novel, intelligent DoS attacks against C-V2X: targeted sidelink jamming and sidelink resource exhaustion. Our attacks demonstrate different ways an intelligent adversary can dramatically degrade the availability of C-V2X for one or many vehicles, increasing the likelihood of fatal vehicle collisions. Through hardware experiments with software-defined radios (SDRs) and state-of-the-art C-V2X devices in combination with extensive MATLAB simulation, we demonstrate the viability and effectiveness of our attacks. We show that targeted sidelink jamming can reduce a targeted vehicle's packet delivery ratio by 90% in a matter of seconds, while sidelink resource exhaustion can reduce C-V2X channel throughput by up to 50% in similarly short order. We further provide and validate detection techniques for each attack based on cluster and regression analysis techniques and propose promising, preliminary approaches to mitigate the underlying vulnerabilities that we expose in the PHY/MAC layers of C-V2X.

Contents

List of Figures

1 Introduction

The emerging family of technologies known as Vehicle-to-Everything (V2X) communication promises a wide variety of societal benefits that range from the realization of fully autonomous vehicles to the elimination of most traffic gridlock. V2X comprises several communication technologies including vehicle-to-vehicle (V2V), vehicle-to-infrastructure (V2I), vehicle-to-pedestrian (V2P) and others [1]. Of these, V2V is considered among the most beneficial due to its anticipated benefits for improving roadway safety. In V2V communication, so-called smart vehicles are able to wirelessly communicate directly with each other, allowing every vehicle to be aware of the location and motion of other vehicles in the area. Thus, V2V has the particular and unique benefit of facilitating crash avoidance in non-line-of-sight (NLOS) situations where a driver or onboard sensors (e.g., cameras, LiDAR) would not be able to perceive or give warning of an imminent collision. One example of such a situation, at a "blind corner" intersection, is given in Figure 1. In the U.S. alone, V2V is projected to reduce the number of vehicle crashes and collisions that occur each year by up to $600,000$ [2], potentially saving thousands of lives annually.

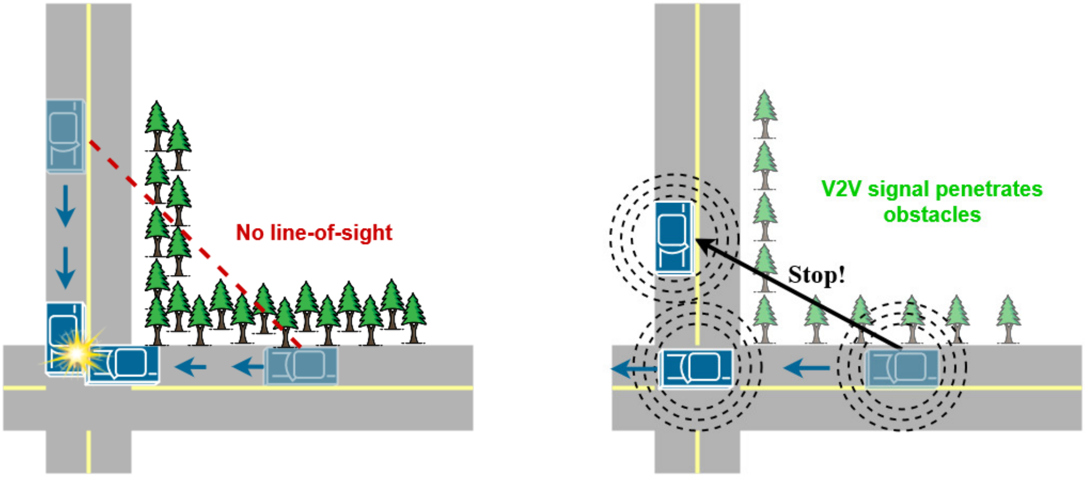

(a) No V2V communication.

(b) V2V communication.

Figure 1: An example of how V2V can be used to avoid a collision in a NLOS scenario.

With the potential to revolutionize the transportation sector through a massive reduction in traffic congestion [3]; advancement towards fully autonomous consumer vehicles [4]; and significant reductions in roadway crashes, injuries, and fatalities [2]; V2V is a highly desirable technology. However, precisely due to the safety-critical nature of V2V, it is imperative that proper steps be taken to ensure V2V communication technologies are secure before human lives are entrusted to them. Drivers—and, eventually, autonomous vehicles—will heavily depend on warnings and guidance originating from V2V messages to make split-second decisions about when to swerve, brake, or otherwise violently maneuver to avoid an anticipated collision. Therefore, any malicious action that causes the V2V system to become unreliable, untrustworthy, or unavailable could have very serious impacts on the safety of drivers and passengers as well as pedestrians and other bystanders.

As V2V has begun to move from theory to reality over the last several years, two major technologies have competed for dominance in the burgeoning [5] V2V market: Dedicated Short-Range Communication (DSRC), based on a decentralized version of IEEE 802.11 [6], and Cellular Vehicle-to-Everything (C-V2X), based on LTE and 5G cellular technologies [7, 8, 9]. Being over 20 years old [10], DSRC was the sole V2V technology for many years and was therefore able to gain acceptance in spite of serious limitations such as a low data rate, lack of multi-user access, and unacceptably high latency [11]. Partly to address these shortcomings, C-V2X was introduced in 2016 in Release 14 of the 3rd Generation Partnership Project (3GPP) standards for cellular communication [7]. Based on LTE cellular technology, Release 14's "LTE-V2X" V2V protocol has since been updated and improved by Releases 15 (2018 [8]) and 16 (2020 [9]). The essence of LTE-V2X is a unique LTE communication mode called "sidelink" that allows direct communication between devices (i.e., vehicles) in place of traditional uplink and downlink (wherein a device can only communicate with other devices through a base station or similar LTE infrastructure). Sidelink communication is ideal for V2V as it removes any requirement for vehicles to be in range of a base station, a capability whose importance will surely be appreciated by anyone who has spent time attempting to find cell service while driving on rural roads.

LTE-V2X is currently the technology of choice for V2X deployments in China [12], may supplant DSRC in Europe [13], and has recently gained both industry [14, 15] and regulatory [16] support in North America. Very recently, 3GPP also standardized 5G-based "New Radio Vehicle-to-Everything" (NR-V2X) in Release 16 [9]. Rather than a successor to LTE-V2X, NR-V2X is envisioned to be a complementary technology [17], with particular usefulness for applications that require very high data rates such as sharing real-time sensor data or streaming live video feeds between vehicles [17, 18]. The most recent NR-V2X standard, Release 16, lays out this relationship in detail as a part of its physical layer specifications [19], wherein it is stated that LTE-V2X will continue to be an active component of future “5G C-V2X” systems [18]. Specifically, Release 16 asserts that LTE-V2X is to be used for safety-critical messages that must meet precise periodicity and latency requirements to ensure the availability of collision-avoidance functionality. The longevity and relevance of LTE-V2X are further supported by the integral role played by LTE-V2X in the current Release 17 draft standard for the “NR-V2X Phase 3” architecture [20], wherein LTE-V2X continues its crucial Release 16 role handling safety-critical messages. Given the current and continuing criticality of LTE-V2X in cutting-edge, 5G-enabled intelligent transportation systems, as well as the general lack of substantive work on security issues in LTE-V2X, this thesis focuses specifically on outstanding security issues in LTE-V2X.

Despite the preeminence and rising popularity of LTE-V2X, the security community has so far largely overlooked it. Much recent work in the literature continues to be based on DSRC (e.g., [21, 22, 23, 24]), and those works which do examine security issues in C-V2X (e.g., [25, 26, 27, 28]) nearly always focus exclusively on NR-V2X. To date, despite the general lack of security considerations included in the LTE-V2X standards [29, 30], there are only a handful of works that have examined security issues in LTE-V2X, and even fewer have examined security at the lower layers of the protocol. Those works which have examined lower-layer security (e.g., [31, 32, 33, 34]) have often proposed either cryptographic or physical-layer approaches to better secure LTE-V2X; however, these and other works deal almost exclusively with either confidentiality or secrecy concerns. This is problematic, because while ensuring the availability of any system is important, in LTE-V2X it is absolutely imperative due to the safety-critical nature of the technology.

In this thesis, we begin to address this gap in the literature by examining the physical (PHY) and MAC-layer security of LTE-V2X from an availability perspective. We particularly focus on how unique aspects of LTE-V2X—for example, its use of decentralized, autonomous selection of time and frequency resources to use for transmitting messages—make LTE-V2X particularly susceptible to unique forms of denial-of-service (DoS) attacks. DoS attacks are particularly insidious in V2V due to the nature of the technology; as described above, human lives are at stake when vehicles with no LOS are on a collision course, and only the availability of LTE-V2X may be able to prevent a tragic outcome. Yet, due to the lack of existing work in this area, a stealthy attacker might be able to deny one or more vehicles the benefits of LTE-V2X at a critical moment, with a range of potential consequences that extends all the way to resulting in avoidable collisions, injuries, and deaths. Motivated by the need to ensure to the greatest extent possible that this sort of malicious action cannot be undertaken, in this thesis we seek to answer the following important questions:

1. Are the PHY and MAC layers of LTE-V2X sufficiently secure as to prevent effective, intelligent DoS attacks by a stealthy adversary?
2. If effective DoS attacks are found at the PHY/MAC layers of LTE-V2X, in what manner can these attacks be effectively detected despite the attacker's efforts to remain as undetectable as possible?

To answer these questions, in this thesis we devise two novel, intelligent DoS attacks against LTE-V2X which exploit unique features of the LTE-V2X PHY/MAC layers to achieve significant DoS effects while remaining undetectable by common detection techniques. The original contributions of this thesis are as follows:

1. A novel attack, *targeted sidelink jamming*, which exploits the distinctive, slot-based structure of LTE-V2X channels, as well as the uniquely periodic nature of V2V messages, to target and jam messages sent by a specific target vehicle with a success rate of up to 93%.
2. A second novel attack, *sidelink resource exhaustion*, in which knowledge of the LTE-V2X MAC-layer is abused by making valid transmissions that cause severe degradation of LTE-V2X channel throughput by deceiving other vehicles into using less bandwidth than is actually available, causing a decrease in channel throughput of up to 50%.
3. We experimentally validate each of the above attacks using SDRs to successfully attack state-of-the-art commercial LTE-V2X devices, affirming the real-world viability of our attacks.

4. We further provide, through MATLAB simulation, evidence of our attacks' effectiveness in an LTE-V2X channel under varying states of busyness.

5. We evaluate the stealthiness of our attacks using MATLAB simulation, showing that state-of-the-art techniques for detecting V2V DoS attacks are (at best) of limited use for detecting our novel attacks.

6. We propose superior, demonstrably more effective techniques for detecting our attacks.

7. Finally, we propose preliminary mitigation approaches to address the vulnerabilities in the PHY/MAC layers of LTE-V2X which we exploited with our attacks.

This thesis is structured as follows. In Section 2, technical information about the lower-layer structure and behavior of LTE-V2X communication is provided along with a primer on V2V communications. Section 3 provides a focused overview of related work on lower-layer security in wireless systems generally and LTE-V2X in particular. Section 4 introduces our contributions and provides some common assumptions for our threat models, while Section 5 introduces our targeted sidelink jamming attack. Section 6 presents our sidelink resource exhaustion attack, following which we compare our two attacks based on a variety of effort and efficiency metrics in Section 7. Finally, we conclude with closing remarks and thoughts on future work in Section 8.

2 Preliminaries

Before describing related work or presenting our contributions, it is necessary to establish some background on LTE-V2X communication as well as relevant V2V fundamentals.

2.1 LTE Vehicle-to-Everything (LTE-V2X)

LTE-V2X is based on LTE-Sidelink (“sidelink”), which was introduced in 3GPP Release 14 [7] as a V2X-specific improvement on LTE Device-to-Device (D2D), an antecedent technology introduced in 3GPP Release 12 [35] for public safety use (e.g., to allow direct communication between firefighters’ cell phones outside of LTE network coverage).

2.1.1 LTE Sidelink Mode 4

Specifically for V2V, Release 14 introduced sidelink Mode 4, the only sidelink mode which allows direct, device-to-device communication in situations where neither device is “in-coverage” (i.e., in communication range of an LTE base station). Mode 4 *uniquely* does not require base stations (or any other LTE infrastructure) to synchronize devices in time and frequency or coordinate resource allocation among user equipments (UEs). Instead, UEs operating under Mode 4 use global navigation satellite systems (GNSS) (e.g., GPS) for synchronization in time. Using GNSS for time synchronization replaces the use of primary and secondary synchronization signals in traditional LTE; consequently, Mode 4 UEs consider all LTE subframes to be millisecond-aligned on GNSS time. In this manner, each $1\,ms$ LTE subframe (see Section 2.1.2) becomes a “time slot” in which one or more transmissions can occur. Although this is critical for ensuring the ability of each vehicle to receive messages without needing to synchronize individually with each transmitting unit, this use of GNSS as a global time reference makes it very easy - far easier than in traditional LTE - for an attacker to accurately anticipate and act against transmissions made in the sidelink channel. This can be a significant security problem, as we explore in Section 5. Finally, it is important to establish that UEs operating in sidelink Mode 4 are required to use their maximum transmit power of $23\,$dBm at all times [36]. This requirement impacts our threat model and assumptions (Section 4) as well as the impacts of our attacks (Section 7).

2.1.2 Structure of the LTE-V2X PHY layer

LTE sidelink supports either 10 or $20\,$MHz channels. This thesis utilizes the $10\,$MHz configuration, a common assumption in the literature. In the time domain, LTE-V2X uses slot-based scheduling. Within each $10\,ms$ LTE frame, ten $1\,ms$ subframes form the time slots used for transmitting or receiving signals. Each LTE frame is globally synchronized in time but locally identified by sequential system frame numbers (SFN) that allow a device to track frames as they elapse over time. The subframes within each frame are in turn identified by a "sidelink frame index" (SFI) between $0-9$. This is ideally suited to the highly periodic traffic used in V2V as it allows easy calculation of transmission times (e.g., $100\,ms$ periodic transmissions might go out in SFNs $2, 12, 22, ...$ using SLI 6 in each of those frames).

In the frequency domain, the $10\,$MHz channel comprising 50 LTE resource blocks (RBs) is divided into five $2\,$MHz subchannels with 10 RBs each. Each subchannel is further divided into control and shared (i.e., data) channels, with the first two RBs of the subchannel allocated to the sidelink control channel (PSCCH) and the remainder to the sidelink shared channel (PSSCH). The time- and frequency-domain structure of LTE-V2X is illustrated in Figure 2. As indicated in Figure 2, a sidelink transmission can use one or more subchannels within a subframe, but any single transmission must occur entirely within one subframe. A transmission is broken into a sidelink control information (SCI) message, which is transmitted in the PSCCH, and a transport block (TB), which carries the message data in the PSSCH. Importantly, no TB can be decoded without first decoding the associated SCI message; thus, as exploited by the attack in Section 5, the dependence of the data channel on the control channel can be a significant security problem in LTE sidelink.

2.1.3 Semi-persistent scheduling (SPS)

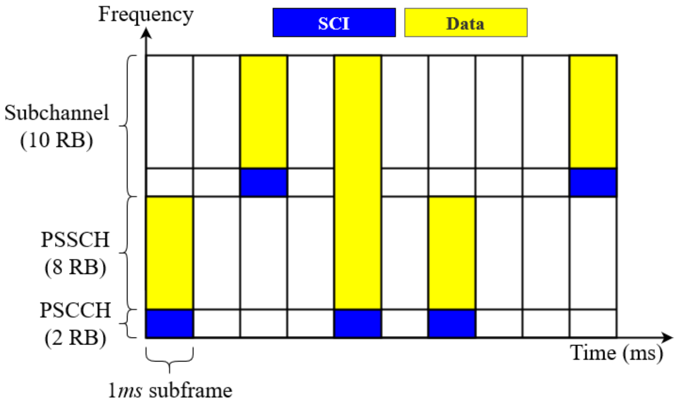

Figure 2: LTE sidelink frame structure.

In the absence of LTE infrastructure to centralize resource scheduling, UEs operating in sidelink Mode 4 use a sensing-based scheduling algorithm called *semi-persistent scheduling*, or SPS [19]. SPS is designed to minimize packet collisions without requiring vehicles to directly coordinate with each other, without the performance-inhibiting "backoff" mechanisms of systems like 802.11, and to make use of the uniquely (within LTE) periodic nature of V2V communication. For example, V2V basic safety messages are sent periodically by every vehicle at a rate of 10 Hz (see Section 2.2). Since all vehicles transmit with the same periodicity, it is easy for a vehicle joining the system to listen to the channel, determine which radio resources (i.e., which subchannel(s) in which subframes) are not being used, and choose unused radio resources to start using for its own transmissions.

SPS also requires vehicles to periodically choose new radio resources (i.e., subframes and subchannels) through "resource reselection" at regular intervals. This requirement is due to the highly mobile nature of V2V. If two or more vehicles are using the same radio resources and come within communication range of each other, then their packets will start colliding and no messages will get through from either vehicle. Forcing resource reselection at regular intervals ensures that if and when situations like this occur, the effect will not last more than several messages (at most). When messages are sent at a rate of 10 Hz, SPS requires resource reselection after every c transmissions, where $c \in [5, 15]$. The value of c is randomly chosen from this interval after each resource reselection and decrements by one after each transmission, triggering the next resource reselection when $c = 0$. Importantly, vehicles can choose to continue using their current resources as well; each time resource reselection occurs, a vehicle will decide whether to choose new resources with pre-configured probability $P \in \{0, 0.2, 0.4, 0.6, 0.8\}$. The effects of different choices for P are explored in Section 5.4.

The resource reselection process is relatively straightforward. For a 1000-subframe ($1\,s$) "listening period" prior to resource reselection, a vehicle will listen to the LTE-V2X channel and create a set of "candidate radio resources" (CRR) from which it ultimately selects new resources to use. For a message periodicity of $100\,ms$, a vehicle needs to select a system frame index, a subframe index within those frames, and a base subchannel within that subframe. For example, the resource set $(3, 2, 4)$ would mean the vehicle transmits in the fourth subchannel of the second subframe of the third out of every ten LTE frames. At reselection time (immediately after the listening period), the vehicle reviews the data collected during the listening period and reduces CRR by excluding resources which meet all of the following criteria:

1. A valid SCI message was received in the sidelink control channel
2. A valid data transport block (TB) was received in the sidelink shared channel using the subchannel(s) indicated by the SCI message
3. The average "reference signal received power"[1] for the subframe and subchannel(s) used for the TB exceeds a defined threshold TH_{rx}

After excluding all resources that meet these criteria, the vehicle checks whether the number of resources left in CRR comprise at least 20% of the total resources in one sensing window. If this is not the case, the vehicle increases TH_{rx} by 3 dB and repeats the process. This allows the vehicle to select resources which were sensed to be in use but had a weak signal, and thus most likely belonged to vehicles which are a large distance away and/or moving away from the receiver.

[1] Defined in 3GPP TS 36.213 [19]; essentially just an average signal strength measurement for the demodulation reference symbols of the shared channel.

2.2 Fundamentals of V2V Communication

Although this thesis is primarily focused on LTE-V2X communication itself, some principles of general V2V communication are important to establish. Since the primary use for V2V is collision avoidance, one of the most important V2V messages is the basic safety message, or BSM. V2V-equipped vehicles broadcast BSMs periodically at intervals between $20-100\,ms$. This thesis uses $100\,ms$ as a standard interval, as is common in the literature, but it is important to note that the $20\,ms$ interval is allowable (as this is a key part of the attack in Section 6). BSMs contain basic information about a vehicle's location and motion (e.g., GPS position, speed and direction of travel) that allows receiving vehicles to be aware of and, if necessary, take action to avoid a collision with the sending vehicle. Under application-layer standards (e.g., [37]), BSMs also carry potentially identifying information about a vehicle like its color, dimensions, make/model, and more. As the specific data reported in a BSM may change, the size of BSMs may vary over time. However, based on experiments with commercial LTE-V2X devices, we assume throughout this thesis that a standard BSM occupies 2 LTE-V2X subchannels within a subframe.

It is important to note that while this thesis is focused on security issues at the lower layers of LTE-V2X, there are upper-layer security requirements in place that are intended to mitigate certain types of attacks. These security requirements, which are general to V2V and do not vary between LTE-V2X and other technologies, are defined primarily in IEEE 1609.2-2016 [38], its amendments [39, 40], and its subordinate standard IEEE 1609.2.1-2020 [41]. The 1609.2 suite includes such security requirements as mandating digital signatures to authenticate BSMs, using public-key certificates to validate message signatures, requiring that all authentication be made pseudonymously to protect privacy, and so on. These upper-layer security requirements are worth noting in order to (a) acknowledge that there are security considerations made in other areas of the V2V protocol stack, and (b) to contextualize some comments referring to these standards made further on in this document. Also, on rare occasions it is possible—though not, to the current body of knowledge, in LTE-V2X—for security mechanisms at the upper layers to be compromised by lower-layer security concerns. The interested reader is encouraged to refer to [42] for details on one such instance.

3 Related Work

Lower-layer security in wireless communications is hardly a new concern. In fact, physical layer security problems—for example, eavesdropping on sensitive communications—date back at least as far as the European battlefields of 1914 [43]. Other physical layer threats, including jamming attacks and more sophisticated forms of eavesdropping, emerged largely out of World War 2 [44]. However, physical layer security in the modern era is no longer exclusively a concern of the battlefield. Due to the everyday ubiquity of wireless communication in modern life, as well as the development of wireless systems more sophisticated than could have been imagined in those early days when speaking in an obscure foreign language was enough to "secure" one's sensitive wireless traffic [45], the security of wireless systems is now a concern for all members of society. This section begins with an overview of existing work on general DoS threats in V2V. We then provide a more detailed review of the few existing works which directly relate to the focus of this thesis—DoS attacks against LTE-V2X—followed by an overview of selected works that establish the state-of-the-art techniques for detecting and mitigating such attacks.

3.1 Denial-of-Service Attacks Against V2V

At the PHY/MAC layers, the most common form of DoS attack in V2V is jamming. Jamming attacks encompass a broad range of threats to V2V communication, but Benslimane and Nguyen-Minh [46] provide a useful starting point for defining and discussing V2V jamming. In particular, they choose packet delivery ratio (PDR) as the metric by which a jamming attack can be evaluated (and, potentially, detected). PDR is also the primary evaluation metric for the attacks we present in Sections 5 and 6. The authors of [46] point out that certain types of jamming attacks, particular periodic jamming attacks, are generally more successful in V2V than in other technologies due to the periodic, predictable nature of V2V transmissions. We exploit this V2V weakness through our first attack in Section 5.

Variations of V2V jamming were examined by Puñal *et al.* [47] in 2015. Their work describes three types of V2V jammers: constant jammers, which transmit continuously on some defined frequency band(s), periodic jammers, which alternate between transmitting and not transmitting in defined intervals, and reactive jammers, which transmit in response to sensing energy “above a certain threshold” [47] on a channel. These classifications are commonly accepted in the literature and so are important to mention, although the two attacks presented in this paper generally fall outside these definitions. Our targeted sidelink jamming attack (see Section 5) is similar to a reactive jammer in that its parameters (e.g., transmission subframe) are defined by a target vehicle’s transmissions, but the jamming pulses themselves are predicted based on the first received message rather than based on each individual transmission by the target. Our other attack, based on resource exhaustion rather than directly jamming transmissions, is most similar to a periodic jammer within the paradigm of [47].

In the literature, jamming attacks against V2V take diverse approaches. Hussein, Mohamed and Krings [48] present a hybrid jamming attack against DSRC which combines reactive jamming and periodic jamming to suppress the periodic transmission of BSMs. Under DSRC, unlike in LTE-V2X, only one vehicle is permitted to transmit at a time. Thus, the authors were able to design a system which sensed when a vehicle was about to transmit and reacted by sending periodic bursts of energy to keep the medium busy with varying-length jamming intervals, causing BSMs that were supposed to be periodic to be queued while waiting for the medium to be free and creating a backlog of messages at the transmitter. This attack specifically exploited the medium contention mechanism of DSRC, so their approach is not applicable to LTE-V2X, but this work illustrates one approach to jamming V2V communications. In a similar but different vein, in prior work [42] we developed a reactive jamming attack against DSRC which exploited the nature of 802.11 transmission and the predictable structure of V2V BSMs to reactively jam messages from a specific target vehicle in real time. We used careful, deliberate decoding of incoming DSRC signals to recover identifying fields and identify the sender of each incoming message before the entire message was received, allowing us to rapidly send a reactive jamming pulse to corrupt the remainder of the message as it arrived at other receivers. This type of attack is highly efficient and effective, but it cannot be applied against LTE-V2X due to the differences in PHY-layer techniques between the two technologies. In DSRC, since only one vehicle can transmit at a time, each incoming message can be processed sequentially. LTE-V2X is very different; since multiple vehicles

Table 1: Related works organized by V2V technology and primary area of focus.

Related Work	**Technology**	**Primary Security Focus**
Alipour-Fanid, Dabaghchian and Zeng [22]	DSRC	Availability
Benslimane and Nguyen-Minh [46]	DSRC	Availability
Feng and Haykin [49]	DSRC	Confidentiality
Gu *et al.* [50]	DSRC	Availability
Hussein, Mohamed, and Krings [48]	DSRC	Availability
Lyamin *et al.* [21]	DSRC	Availability
Pirayesh *et al.* [24]	DSRC	Confidentiality/Integrity
Puñal *et al.* [47]	DSRC	Confidentiality
Sun *et al.* [23]	DSRC	Availability
Twardokus *et al.* [42]	DSRC	Confidentiality
Lautenbach *et al.* [25]	NR-V2X	Confidentiality
Lai *et al.* [28]	NR-V2X	Availability
Lu *et al.* [26]	NR-V2X	Confidentiality
Nguyen, Lin and Hwang. [27]	NR-V2X	Availability
ElHalawany, El-Banna and Wu [31]	LTE-V2X	Confidentiality
Li *et al.* [51]	LTE-V2X	Availability
Liu *et al.* [33]	LTE-V2X	Integrity
Luo *et al.* [34]	LTE-V2X	Availability
Trkulja, Starobinski and Berry [52]	LTE-V2X	Availability
Wang *et al.* [32]	LTE-V2X	Integrity
Our work	**LTE-V2X**	**Availability**

may transmit at the same time, a receiver collects one subframe ($1\,ms$) of samples at a time, then searches through the received data to iteratively decode any and all messages that were received in that subframe of data. Thus, in LTE-V2X a message can only be recovered after it is received in its entirety, precluding use of our technique from [42] to attack LTE-V2X communication. Both [42] and [48] demonstrate some of the many reasons why most existing work on V2V DoS attacks, which are usually designed for DSRC, cannot be applied to attacking LTE-V2X.

3.2 Denial-of-Service Attacks Against LTE-V2X

To the best of our knowledge, there are only two existing works that describe lower-layer DoS attacks specifically designed for use against LTE-V2X. The first of these (Li *et al.* [51]) primarily focuses on detecting such an attack and is described in Section 3.3. In the other work, Trkulja, Starobinski and Berry [52] seek to demonstrate through modeling and simulation that LTE-V2X is prone to intelligent DoS attacks; specifically, they explore a process they call "adversarial resource block selection" in which one or more vehicles contort their use of SPS (see Section 2.1.3) to probabilistically choose resources that are *more likely* to be in use by other vehicles rather than less likely (as SPS is intended to do). This approach is somewhat similar to the motivation for our attack in Section 6, as like us their objective is to abuse SPS for malicious ends. However, where the authors of [52] particularly focus on collusion attacks involving multiple malicious vehicles working in tandem, our attacks are intended to be executed by a single attacker. The system model in [52] is also somewhat oversimplified; for example, instead of subchannels and subframes they simply treat each transmission as using one "resource block" and give little attention to the granular details of LTE-V2X's slot structure and radio resource allocation procedures. Some details of the system model in [52] also do not align with LTE-V2X standards (e.g., they use a channel with 200 resource blocks, which is not possible in LTE-V2X as the maximum $20\,\mathrm{MHz}$ bandwidth equates to 100 resource blocks). While our simulated channel models are also simplified compared with real-world operating conditions, our model accounts for significantly more details of SPS and the PHY/MAC structure of LTE-V2X than [52] and we have taken care to ensure the parameters we do model in our simulations are based on applicable LTE-V2X requirements from 3GPP (e.g., [7, 19]).

3.3 Detecting V2V Denial-of-Service Attacks

The most common approaches to detecting DoS attacks in V2V are based on evaluating PDR for the overall system. PDR can have different definitions in different contexts. In this thesis, we define PDR as a ratio of the number of packets transmitted by a vehicle (or all vehicles, for a system) to the number of packets received correctly by at least one vehicle. This often involves developing a method of estimating the number of periodic messages (e.g., BSMs) that should be received in a given period of time, and then developing a threshold to raise an attack alert when PDR falls too far below expectations. Benslimane and Nguyen-Minh present one such system in [46]. They take a MAC rather than PHY-layer approach to detecting jamming, differentiating [46] from much related work. The central tenet of their detection scheme is a probabilistic system of classifying time slots where no BSMs are received as either acceptable, the result of a packet collision due to normal operation, the result of interference or noise due to unintentional activity, or as the result of malicious jamming. While their scheme has great success in certain scenarios when the number of vehicles in a system is restricted and known *a priori*, they acknowledge that their detection mechanism is only around 50% accurate (at best) in less controlled scenarios. This is a common shortcoming of existing work on V2V DoS detection as many such works (e.g., [21, 22, 53]) rely on *a priori* knowledge of the number of vehicles within communication range. While we evaluated our detection techniques using simulations with defined numbers of vehicles, our detection techniques do not require advance knowledge of the number of vehicles in the system to be effective (although their effectiveness, as discussed in Sections 5.6 and 6.7, does vary with the number of vehicles in the system).

Li *et al.* [51] proposed a "resource exhaustion" attack against LTE-V2X in which an attacker floods the network with high-priority packets that require immediate attention, leaving no resources left to handle ordinary (i.e., lower priority) traffic from other vehicles. The motivation for their attack is strikingly similar to ours in that the goal is for an attacker to exhaust network resources by making legitimate transmissions; however, there are critical differences between our attack in Section 6 and that of [51]. First, their attack is against LTE Sidelink Mode 3, which involves vehicles communicating to each other through an LTE base station (in a similar manner to mobile phones). In fact, the crux of their attack is flooding a base station with high-priority requests, an approach which is entirely inapplicable to the our Mode 4 scenarios. Second, the authors of [51] straddle the border between a lower- and upper-layer attack by relying on packet priority, which may be technically a MAC-layer element but is often set and processed by the upper layers. Our attack, which targets the MAC-layer scheduling algorithm of LTE-V2X in Mode 4 (see Section 2.1.3), is an entirely lower-layer approach that is quite distinct from this existing work. In terms of effectiveness, the authors of [51] show simulation results indicating that their attack can achieve a more significant reduction in channel throughput than ours, in some cases up to 100%. However, they achieve these results at the expense of detectability; as their attack must actively request resources from a base station—which requires authentication [38]—the attacker must identify herself to the network in order to execute the attack. Even if the attacker forges or steal credentials, her transmissions may still easily be identified by her exclusive use of the resources requested from the base station, once her requests are identified as malicious (by, for example, using the scheme devised in [51]). Thus, while effective, the attack presented in [51] is not an intelligent DoS attack like ours.

3.4 Mitigating V2V Denial-of-Service Attacks

Many works have examined the problem of mitigating V2V DoS attacks at all layers of the protocol stack. Some, such as [54] and [55], have discussed DoS attacks at the upper layers and proposed some mitigation techniques which are indirectly of interest but not generally applicable to discussion of mitigating lower-layer DoS attacks in V2V. Of greater relevance to this thesis are those works which have examined mitigating jamming attacks against V2V systems. Feng and Haykin [49] propose one such approach based on application of cognitive risk control (CRC) techniques. Their proposal is to allow vehicles to adapt their behavior, particularly channel selection and transmit power, based on behavioral observation of a smart jammer. Although they provide some promising preliminary results, this work focuses on a limited case where a jammer is attempting to disrupt communication between two specific vehicles. Their approach requires a vehicle to be capable of accurately determining both the distance between itself and the jammer and the distance between itself and the vehicle(s) to whom jammed messages are being sent. It may be possible to fulfill this requirement when considering only two vehicles—though the authors admit this is already challenging—but much less so in situations like ours, where BSMs are being sent to every vehicle within several hundred meters of the sender. Therefore, it is doubtful that the mitigation technique proposed in [49] could be extended to mitigate BSM jamming attacks.

Many mitigation techniques have been proposed that are more suitable for preventing BSM jamming; however, these proposals often suffer from being designed specifically for DSRC; thus, in one way or another, they are often inapplicable to LTE-V2X systems. For example, Gu *et al.* [50] proposed a system to mitigate control-channel jamming in V2V systems. The authors of this work examine the case where an attacker jams the V2V control channel to disrupt the delivery of V2V messages to vehicles within range of the jamming signal, a conceptually similar attack to the one we present in Section 5. Their idea is to use cooperative relaying of control-channel messages—i.e., vehicles for whom the control channel is jammed can still receive control-channel messages via peer "cooperative relay" vehicles unaffected by the jamming—to mitigate the risk of a control-channel jammer crippling an entire V2V system. This approach is efficient and effective; however, it is designed specifically for the particular channel arrangement of DSRC, which differs fundamentally from that of LTE-V2X. In DSRC, a given frequency band (e.g., $5.85-5.92$ GHz) is subdivided into several 10 MHz channels, one of which becomes the control channel while the others become shared channels for transmitting data like BSMs. Importantly, this means that the control channel can be jammed *without* directly impacting the shared channels. This is impossible in LTE-V2X, which requires every shared-channel message (e.g., BSM) to have an associated control-channel message informing the receiver of how to decode the transmitted data (see Section 2.1.2). Also, while DSRC only has one control channel, LTE-V2X has one for every subchannel, making any sort of message relaying technique difficult to apply as a relay vehicle would need to monitor at least five different control channels at once for messages to relay. Thus, while a promising approach for mitigating control-channel jamming in DSRC, the technique of [50] cannot be effectively used to mitigate control-channel jamming attacks in LTE-V2X systems. Pirayesh *et al.* [24] proposed another mitigation technique designed for DSRC that does not work for LTE-V2X. In [24], the authors investigated mitigation of high-power jamming attacks against V2V communication. Their work proposes two adaptations to the receiver design in IEEE 802.11p (i.e., DSRC) which they show allow for mitigation of jamming attacks with a jamming-to-signal ratio as high as 25 dB. The authors provide substantial experimental validation of their approach, which is based on spatial filtering using a MIMO receiver; however, because the proposed modifications they make for DSRC depend heavily on the existence of training fields in V2V message preambles—neither of which exist in LTE-V2X—this technique is also not applicable to LTE-V2X. There are many more examples of this shortcoming of existing works, but [50] and [24] provide illustrative examples of why DSRC-based mitigation approaches, which are by far the most common in the literature, cannot be applied to mitigating jamming attacks in LTE-V2X systems.

Given the lack of mitigation techniques for LTE-V2X jamming in the literature, it is useful to look at existing work that has examined related problems (e.g., latency and throughput issues, reducing packet loss) from a system performance perspective. For example, Mughal *et* *al.* [56] proposed modifying the sending and receiving procedures of LTE-V2X in response to the problem of control-channel contention leading to high levels of packet loss when SCI messages collide with each other and render associated V2V messages (e.g., BSMs) irrecoverable. While their work is motivated by a desire to decrease the packet error rate that occurs in LTE-V2X due to SPS resource reselection conflicts between vehicles, their approach is also useful as a technique to mitigate the attack we propose in Section 5. In [56], the authors suggest changing the arrangement of control and data channels within LTE-V2X transmissions so that every data transmission (e.g., BSM) does not require successfully receiving and decoding an SCI message as a prerequisite for retrieving the transmitted data. Their scheme involves embedding control information within data transmissions rather than sending control and data information separately, alleviating the problem of congestion in (and, indeliberately, certain attacks against) the LTE-V2X control channel. We discuss this work further in proposing mitigation techniques for one of our attacks in Section 5.7. Another work which identifies techniques for performance improvements that can be co-opted for mitigating attacks is [57]. In that work, Nabil *et al.* examine performance problems (e.g., low PDR at high vehicle density) with the SPS algorithm. They particularly examine the effect that the resource reservation interval (i.e., the period between resource reselections) has on overall system performance. Among other conclusions, they find that increasing the resource reservation interval increases PDR by lowering the probability that vehicles select conflicting resources during SPS resource reselection, though only with certain costs in terms of data rate and latency. This is an interesting approach to consider with respect to the attack we propose in Section 6, and we discuss it further in our proposed mitigation techniques (see Section 6.8).

4 Novel Denial-of-Service Attacks Against LTE-V2X

Having established the necessary technical background and reviewed related works, the remainder of this thesis is focused on answering our two research questions. In Sections 5 and 6 we present two novel DoS attacks, each of which exploits different fundamental flaws in the PHY and MAC layers of LTE-V2X, thereby demonstrating that the lower layers of LTE-V2X are not sufficiently secure as to prevent intelligent, protocol-aware DoS attacks from being effective. We also devise and validate methods for detecting each attack and propose preliminary ideas for mitigating them, providing a path to securing LTE-V2X against the threats we have identified.

In presenting each attack, we use a common parlance and make some common assumptions about both the attacker and the LTE-V2X environment. We assume a single attacker, "Eve," is to undertake each of our attacks. Eve is generally assumed to have similar capabilities to an ordinary vehicle that is equipped with LTE-V2X technology. So, Eve is unrestricted in her movement and may be either mobile or stationary. She has the ability to receive and transmit signals on channels in the 5.9 GHz LTE-V2X frequency band; further, she is able to transmit V2V messages which are syntactically valid (i.e., compliant with communication requirements like message size) as well as semantically valid (i.e., contain meaningful and appropriately formatted data) as necessary. We assume Eve's communication device may be equipped with multiple antennas, though they may not always be necessary. Finally, to ensure Eve cannot be easily identified as malicious due to non-compliance with basic operating requirements, we set Eve's transmit power to 23 dBm per [36] and we limit her to transmitting no more than one message every $20\,ms$ per [19].

5 Targeted Sidelink Jamming Attack

In our first attack, *targeted sidelink jamming*, the attacker Eve aims to deny the collision-avoidance benefits of V2V to a single, specific victim vehicle ("Alice"). Eve pursues this objective by attempting to directly jam as many of Alice's BSMs as possible, thereby decreasing Alice's safety on the road as other vehicles receive fewer of her BSMs, consequently become less aware of her movements, and finally become more likely to maneuver in an unsafe manner that may result in a collision with Alice. One example of such a scenario is depicted in Figure 3.

The informed reader may immediately ask how Eve is able to identify when a BSM has come from her specific target, Alice, when V2V security standards (i.e., [38, 41]) specifically require anonymization of BSM contents. A complete answer to this question is beyond the scope of this thesis; however, there are several ways this identification might be accomplished. It is important to note that the contents of a BSM, as defined by the industry standard [37], include such potentially identifying information as the make, model, length, width, color, etc. of the vehicle which sent the BSM. In certain contexts, this alone may be sufficient; for example, if Alice has a particularly unusual combination of make, model, and color, then her messages may be easily identified by these aspects. Of course, it is more likely that Alice is not visually unique, so other approaches might be necessary. If Eve is able to maintain even intermittent visual contact with Alice, which may be likely if Eve is following Alice in a separate vehicle, then such technologies and techniques as directional antennas, angle-of-arrival analysis, etc. may be used to isolate Alice's BSMs from others. These techniques have been proposed for verifying positional and motion claims of BSMs (e.g., [58, 59]) and might be co-opted for Eve's nefarious intentions. In any case, these are a small subset of the potential methods Eve might use to identify Alice's BSMs, and a further discussion is left to future work. From here on, Eve's ability to identify Alice's BSMs is assumed.

5.1 Attack Procedure

Eve begins by listening to the channel until she receives a BSM from Alice. Assuming, as we do, that Alice is complying with LTE-V2X standards, this will occur within $100\,ms$. When a BSM

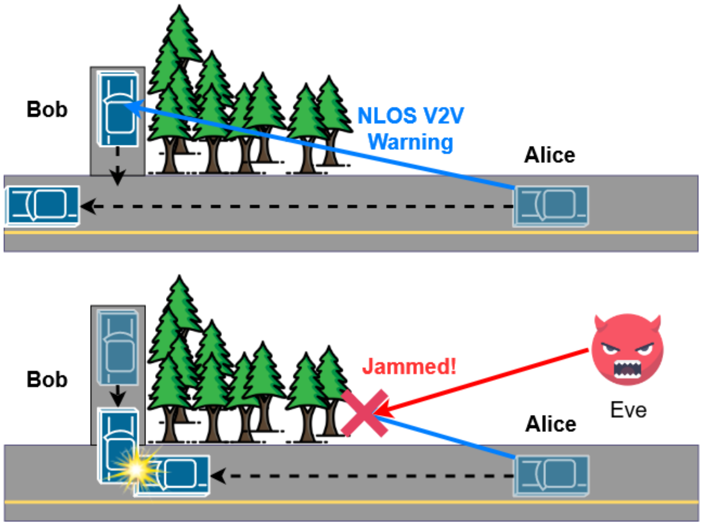

Figure 3: Potential consequences of the targeted sidelink jamming attack.

from Alice arrives, Eve records the SFN, SFI, and the subchannel(s) of that BSM as Alice's. Now, since Eve knows that the standard BSM interval is $100\,ms$, she can accurately predict the SFN and SFI of Alice's next BSM; further, since Alice's subchannel(s) will not change until resource reselection occurs (see Section 2.1.3), Eve also knows which subchannel(s) Alice will use to transmit her BSM. Collectively, this means Eve can predict both the time and frequency resources not only of Alice's next BSM, but of her next several BSMs, with near-perfect accuracy. Put another way, if Alice's first BSM arrives in subframe (SFN, SFI), then Eve knows Alice's future BSMs will arrive (using the same subchannel(s)) in subframes $(SFN + 10i, SFI)$, where $i \in \{0, 1, ..., n\}$ and $n \in [5, 15]$). Figure 4 illustrates this attack procedure.

At this point, Eve knows exactly which time and frequency resources Alice will use for her next several BSMs. Eve's next step is to attempt to prevent other vehicles from receiving those BSMs, which she does by transmitting her own messages, in the same resources, that collide with Alice's BSMs and make them unrecoverable by receivers. However, Eve does not simply transmit a complete BSM, because this would be inefficient[2]. Instead, to jam each BSM, Eve transmits an SCI message in the control channel to collide with the SCI message associated with that BSM.

[2] Also, it might make Eve more detectable. Consider what would happen if Alice stopped transmitting for some reason and Eve, unaware, continued transmitting in what she believes are still Alice's resources. If Eve transmitted complete BSMs, or even just a large amount of noise across the subchannels (formerly) used by Alice, it might be immediately obvious to an observer that an attacker was present.

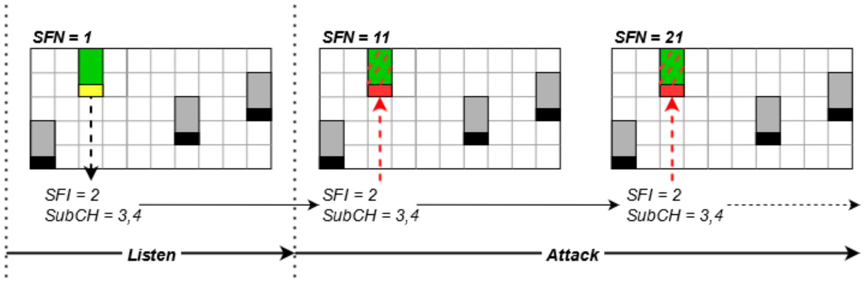

Figure 4: Targeted sidelink jamming attack procedure.

Now, when a receiver attempts to process (SFN, SFI) and recover the BSM, the receiver will be unable to properly recover Alice's SCI message; consequently, the receiver will also be unable to recover the associated BSM[3].

In Figure 4, the SCI message associated with the first of Alice's BSMs to arrive at Eve is shown in yellow and the SCI messages of subsequent BSMs in red, indicating Eve's deliberate collision with them. Incidentally, this approach to jamming Alice's BSMs has the advantage of being highly deniable in LTE-V2X due to its use of SPS, because of which vehicles may, not infrequently, select conflicting resources and transmit colliding SCI messages, an occurrence which appears identically to Eve's actions against Alice. Related work has shown that the overall PDR of an LTE-V2X channel can drop significantly due to this result of using SPS—in extreme cases, by as much as 40% [57]—even without an attacker being present. Therefore, the loss of a small number of BSMs (i.e., the BSMs from Alice which are jammed by Eve) is likely to be attributed to the general, expected packet loss under SPS. We discuss this further in Section 5.5.

Eve continues this pattern of attack, jamming Alice's BSMs at $100\,ms$ intervals indefinitely until it is time for Alice to reselect her resources, which occurs every $5{-}15$ messages with probability $P \in \{0, 0.2, 0.4, 0.6, 0.8\}$ (see Section 2.1.3). Importantly, Eve needs to be able to detect when Alice reselects resources so that Eve can adjust her jamming accordingly. To do this, in between jamming Alice's BSMs, Eve continues to listen to the channel and process incoming BSMs, checking each one to see whether it was sent by Alice. If a BSM is received from Alice using a different SFN, SFI or subchannel than Eve expects, then Eve can infer that Alice has reselected her resources. Eve can then simply update her record of Alice's resource reservation (i.e., SFN,

[3] Recall from Section 2.1.2 that a valid SCI message is required in order to decode any LTE-V2X transmission.

SFI, and subchannel(s)) and continue jamming Alice's BSMs at $100\,ms$ intervals. Thus, with the exception of the first BSM that Alice sends after reselecting resources (i.e., maximally one out of every $5-15$ BSMs), Eve is able to accurately jam all of Alice's BSMs. This observation yields an anticipated degradation of Alice's PDR of between $80\%\,(\frac{4}{5})$ and $93\%\,(\frac{14}{15})$ overall, making this a very effective form of DoS attack.

5.2 Experimental Validation

We evaluated the real-world viability of our attack experimentally, using an SDR and a proof-of-concept software implementation to effectively attack state-of-the-art commercial LTE-V2X equipment.

5.2.1 Experimental hardware

Our experimental setup required a combination of hardware devices. To represent two "vehicles"—the target, Alice, as well as a second vehicle, "Bob," to receive Alice's messages—we used state-of-the-art commercial LTE-V2X equipment from Cohda Wireless [60] (see Figure 5). Testing our attack against these devices, which are widely used in real-world roadway testing of V2V/V2X deployments, strengthens the validity of our results and affirms the real-world viability of our attack. For Eve, we used a USRP B210 SDR equipped with two $5\,$dBi antennas as well as a GPSDO module for GNSS time synchronization. To synchronize all devices in time, we used a LimeSDR to transmit synthesized GNSS signals which were generated using GPS-SDR-SIM [61]. Our experimental setup is shown in Figure 6.

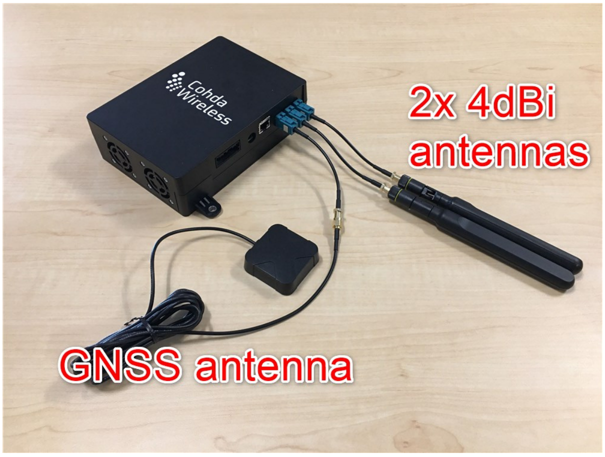

Figure 5: Cohda Wireless MK6C evaluation kit.

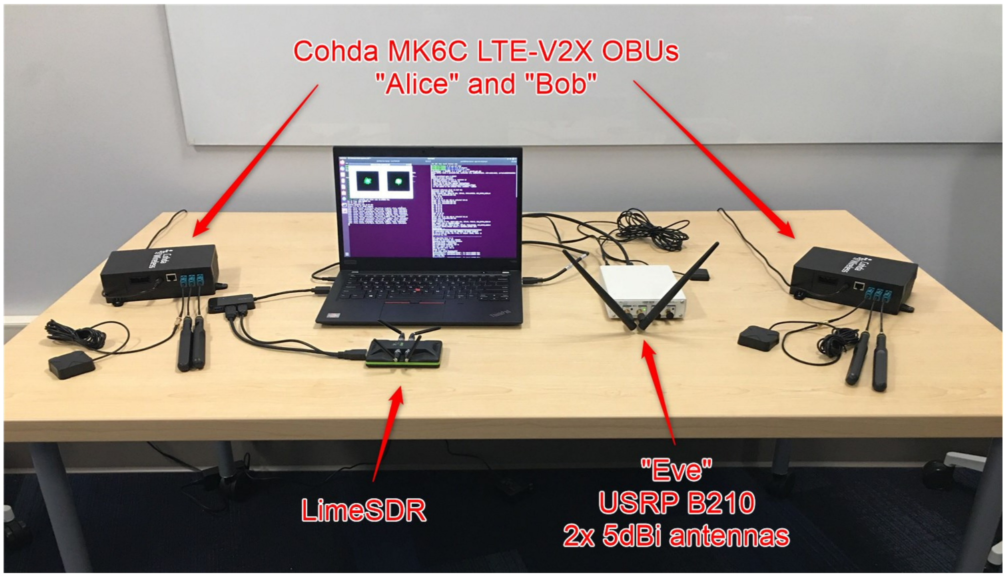

Figure 6: Experimental setup for evaluating the targeted sidelink jamming attack.

5.2.2 Proof-of-concept software implementation

We implemented a proof-of-concept version of our attack in C++. To a limited extent, we made use of open-source libraries from srsRAN [62] as a starting point, and we drew some insight from examination of prior work by Eckermann *et al.* [63]. However, significant extensions to and improvements on these works were required to meet our particular requirements. Among other capabilities, we developed support for real-time analysis of received signals and created the capacity for time-scheduled transmissions of LTE-V2X messages. Further, we undertook the non-trivial task of combining entirely separate transmit and receive functionalities (based on different C/C++ libraries) to develop a sidelink application that can be used with a dual-antenna USRP to emulate a vehicle (or attacker). A part of this sidelink application that was developed in the course of completing this thesis has been submitted for inclusion in an open-source V2V security project we developed in prior work [42], and that component of this thesis work will be publicly released in an upcoming version of that project.

5.2.3 Experimental setup

We configured Alice and Bob (the two Cohda LTE-V2X devices) to transmit BSMs at the standard $10\,\text{Hz}$ rate for a period of ten minutes. To obtain a baseline measurement for PDR in the absence of an attacker, at the conclusion of this period we reviewed the packet capture logs on each device and compared the number of BSMs received by each device to the expected quantity of BSMs sent in a ten-minute period. From this analysis, we determined that the baseline PDR in our controlled experimental environment was nearly perfect ($> 99.85\%$[4]) in the absence of an attacker. This justifies assignment of any more significant decrease in PDR that we observe during evaluation of the attack to Eve's actions rather than to environmental factors or other experimental biases.

To evaluate our attack, we once again configured Alice and Bob to transmit BSMs at a rate of $10\,\text{Hz}$. We set the transmit power for both Alice and Bob to $23\,\text{dBm}$, which is both the maximum allowable and default power setting for the Cohda MK6C devices. We configured Eve, using a dual-antenna SDR as described above, to execute our proof-of-concept implementation of the targeted sidelink jamming attack against Alice. Eve transmits with an approximate[5] power level of $10\,\text{dBm}$ to send her jamming signals against Alice's BSMs. Note that this power level is significantly lower than that used by Alice, an observation of relevance to later discussion about the efficiency of this attack in Section 5.3.

5.2.4 Experimental results

Our experimental evaluation of the attack was very positive. As shown in Figure 7, Eve was able to reduce Alice's PDR below 20% after $1\,s$ (10 BSMs) of jamming. Eve's PDR levels off around 10% over time, confirming the expected $80-93\%$ reduction in Alice's PDR as a result of the attack. One can also see from Figure 7 that Bob's PDR is negligibly affected, which is very important for Eve's ability to remain undetected. These results validate the real-world viability of our targeted sidelink jamming attack against state-of-the-art commercial LTE-V2X equipment. Also, in demonstrating

[4] Alice and Bob received $5,992$ and $5,994$ BSMs, respectively, out of an expected $6,000$ over the course of 10 minutes.

[5] The USRP B210 manufacturer, Ettus Research/NI, does not supply a precise value, specifying only a maximum RF output power of "$>10\,\text{dBm}$" [64]. However, related work has found that $10\,\text{dBm}$ is an accurate estimate for the B210's transmit power at RF frequencies close to $6\,\text{GHz}$ [65].

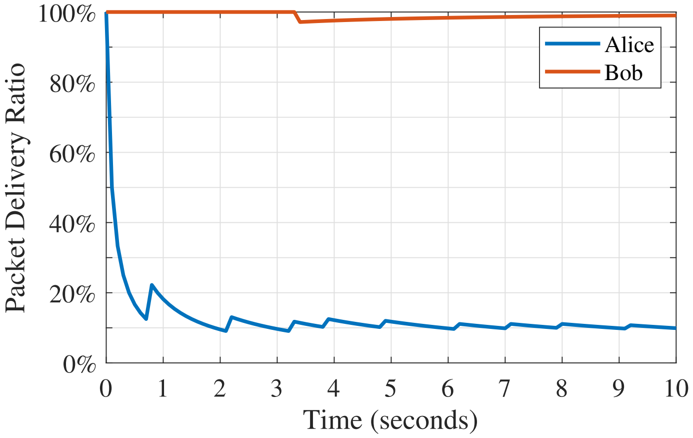

Figure 7: Experimental results showing the effectiveness of targeted sidelink jamming.

the viability of an effective, intelligent, protocol-aware DoS attack against LTE-V2X, we have answered our first research question in the negative. That is, we have demonstrated that the LTE-V2X PHY/MAC layers are *not* sufficiently secure as to prevent such intelligent DoS attacks from being possible.

5.3 Efficiency and Effort Considerations

Beyond demonstrating a proof-of-concept implementation of our attack, some further discussion is in order regarding the effectiveness and considerations required for the attack to be executed in a more realistic environment.

5.3.1 Jamming-to-signal ratio

Jamming-to-signal ratio (JSR), as measured at an arbitrary receiver, is a common way of evaluating the effort that is required for an attacker (i.e., jammer) to successfully prevent a signal from being received. We can calculate the JSR for our experiments using the Friis transmission equation [66] for received power:

$$P_r = \frac{P_t G_t G_r \lambda^2}{(4\pi d_t)^2} \tag{1}$$

where P_r is the power of the received signal, P_t is the transmit power, G_t is the gain of the transmit antenna, G_r is the gain of the receiver antenna, λ is the signal wavelength, and d_t is the distance between transmitter and receiver. From the Friis equation, we can derive a general formula for JSR (substituting subscript J to denote jammer variables or subscript S to indicate those for the sender of the jammed signal, per Table 2) as follows.

$$\frac{J}{S} = JS^{-1} = \left(\frac{P_J G_J G_R \lambda^2}{(4\pi d_J)^2}\right)\left(\frac{(4\pi d_S)^2}{P_S G_S G_R \lambda^2}\right) = \left(\frac{P_J G_J d_S^2}{P_S G_S d_J^2}\right) \tag{2}$$

However, this equation is for a real unit of power (e.g., in Watts), whereas we want to calculate JSR in decibels. JSR in decibels is given by (3) below.

$$\frac{J}{S}\,(\mathrm{dB}) = 10\log\left(\frac{P_J G_J d_S^2}{P_S G_S d_J^2}\right) = P_J + G_J + 20\log(d_S) - P_S - G_S - 20\log(d_J) \tag{3}$$

Finally, using the experimental values given in Table 2, we can use (3) to calculate a JSR that we know from our experiments is sufficient for the attack to effective.

$$\begin{aligned}\frac{J}{S} &= P_J + G_J + 20\log(d_S) - P_S - G_S - 20\log(d_J)\\ &= 10 + 5 + 20\log(2) - 23 - 4 - 20\log(0.5)\\ &= 0.041\,\mathrm{dB}\end{aligned}$$

Given an experimental JSR of $0.041\,\mathrm{dB}$, we know our attack works when the jamming signal has roughly[6] the same strength as the jammed signal. In our parlance, then, Eve's jamming transmissions should *minimally* have the same signal strength as Alice's transmissions at the receiver(s) in order to be certain—insofar as our experimental results can guarantee—that Alice's BSM will be irrecoverable by those receiver(s). This has some obvious ramifications for the real-world execution of our attack. For one thing, Eve will need to consider her positioning with respect to Alice very carefully in order to ensure she can achieve the necessary JSR. Similarly, she may wish to increase the power with which she transmits her jamming signals, potentially as high as the

[6] To be precise, when the jamming signal is $10^{0.041/10} = 1.0095$ times as strong as the jammed signal. However, for the purposes of discussion, this is negligible and can be treated as an effective JSR of 1 ($0\,\mathrm{dB}$).

Table 2: Values for experimental parameters that affect jamming-to-signal ratio.

Experimental Variable	**Symbol**	**Value**
Jammer power	P_J	$10\,\text{dBm}$
Victim's transmit power	P_S	$23\,\text{dBm}$
Jammer antenna gain	G_J	$5\,\text{dBi}$
Victim's transmit antenna gain	G_S	$4\,\text{dBi}$
Distance between jammer and receiver	d_J	$0.5\,\text{m}$
Distance between victim and receiver	d_S	$2\,\text{m}$
Wavelength ($f = 5.92\,\text{GHz}$)	λ	$0.0508\,\text{m}$

limit we set ($23\,\text{dBm}$—see Section 4)[7] where she transmits with equal power to Alice. Figure 8 illustrates an example of these considerations with a more realistic distance of $50\,\text{m}$ between Eve and Alice during the attack. As shown by the left-hand subplots of Figure 8, if Eve uses her maximum transmit power of $23\,\text{dBm}$, she can achieve a JSR of $0\,\text{dB}$ or better over the majority of a $40,000\,\text{m}^2$ area centered on her location. However, if she only uses $10\,\text{dBm}$ transmit power, her effective jamming area drops dramatically (as shown by the right-hand subplots of Figure 8). Of course, a real attacker would take many factors into account when selecting her transmit power and position with respect to Alice (e.g., specific desired outcome, road layout in the area, traffic speed), but the example given in Figure 8 demonstrates the necessity of doing so solely on the basis of achieving the required JSR.

5.3.2 Efficiency

Another important consideration is the level of effort that Eve is required to put in to successfully accomplish her goals (of jamming Alice's BSMs). One way of evaluating this is calculating the *duty cycle* of the attack. Duty cycle has many definitions, but here we define it as the ratio of the jammed bandwidth to the bandwidth that is ultimately affected by the jamming signal. As described in Section 5, each jamming signal is constituted by a single SCI message, which has a width of 2 sidelink resource blocks ($400\,\text{kHz}$). In each instance of jamming, though, an entire BSM is rendered unrecoverable. Since a standard-size BSM uses 20 resource blocks (including the jammed SCI

[7] NB: We showed through our experiments that it is not *necessary* for Eve to transmit with equal power to Alice, though. This gives Eve latitude to vary her transmit power as needed to achieve her goals and does not lock her into using a higher, more detectable power level when doing so is not required.

message), or 4 MHz of bandwidth, the duty cycle for our jammer is $\frac{400e3}{4e6} = 0.1 = 10\%$. This is a respectably low, if not a particularly efficient result. We discuss this further in Section 7.

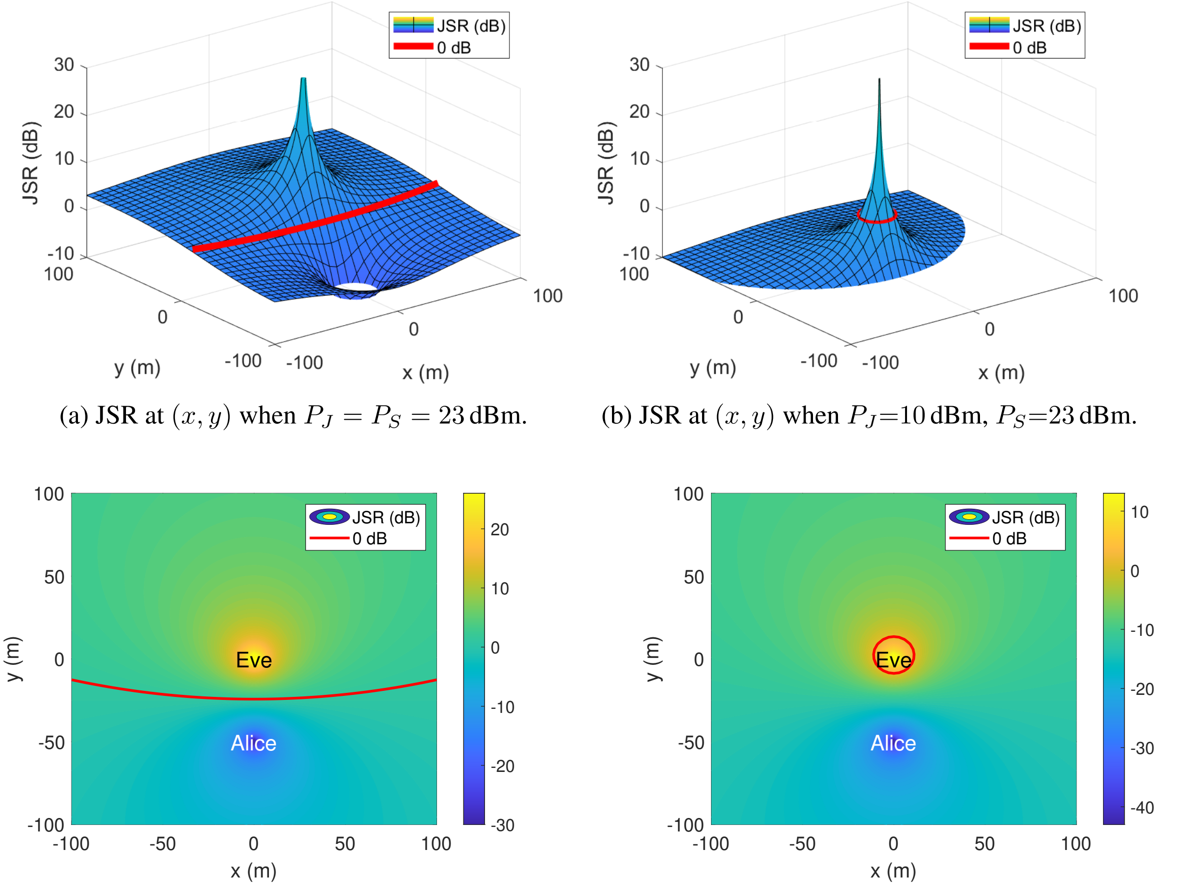

(a) JSR at (x, y) when $P_J = P_S = 23$ dBm.

(b) JSR at (x, y) when P_J=10 dBm, P_S=23 dBm.

(c) Overhead view for $P_J = P_S = 23$ dBm.

(d) Overhead view for P_J=10 dBm, P_S=23 dBm.

Figure 8: JSR for a receiver located at (x, y) within a $40000\,\mathrm{m}^2$ area centered on Eve. Red lines indicate the limits of Eve's effective jamming range.

5.4 Modeling LTE-V2X in MATLAB

Although experimental work with hardware devices was critical to affirm the viability of our attack against real LTE-V2X equipment, we cannot reasonably extrapolate from our experimental results against one device to describe the impact of our attack on an LTE-V2X channel used by many more vehicles. As this evaluation is necessary in order to determine how detectable the attack is, we created an LTE-V2X system model using MATLAB. We modeled a standard 10 MHz LTE-V2X channel with 5 subchannels (as described in Section 2.1.2), using a matrix-based representation of channel resources to track transmissions and record packet collisions. Figure 9 shows an example of how a sidelink frame is represented in our model, with color scaling to indicate which resources were used for transmissions. Following LTE-V2X standards, we configured simulated vehicles to transmit BSMs at the standard $10\,\text{Hz}$ rate in this channel, with each BSM's size set at 2 subchannels (based on commercial standards [37, 67]). Finally, we implemented the SPS algorithm in accordance with 3GPP TS 36.213 [19], configuring each simulated vehicle to regularly reselect resources based on channel usage in the same manner as real vehicles would do.

This channel model is deliberately "perfect"—i.e., the causes for packet loss are limited to packet collisions resulting from SPS and malicious actions by a DoS attacker; equivalently, we assume that every transmitted BSM from any vehicle will be successfully received by every other vehicle *unless* it either collides with another vehicle's BSM (due to SPS resource conflicts) or is blocked by an attacker. Our reason for using this perfect channel model is to create a worst-case scenario for an attacker who wishes to remain undetected. In our model, the attacker's actions must cause sufficiently little effect on overall PDR as to be statistically difficult to distinguish from the packet loss which results from SPS; otherwise, the attack will be detectable. We argue that if the effects of our attack on PDR cannot be reliably distinguished from SPS-based packet loss alone in this perfect channel, then it follows that our attack will be similarly (or even more) difficult to detect using PDR in a realistic channel where factors like noise, fading, etc. will all contribute to cause additional packet loss and make distinguishing the packet loss caused by an attack even more challenging.

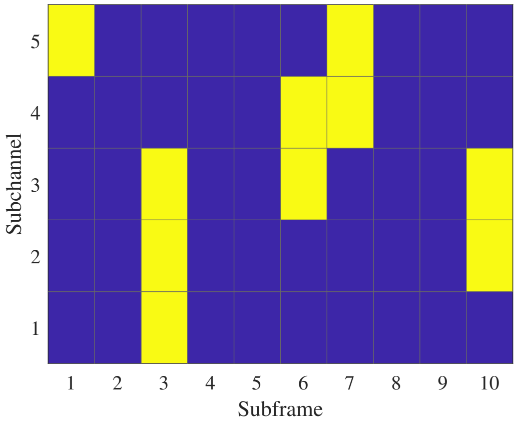

Figure 9: A $10\,ms$ sidelink frame as represented within our LTE-V2X model. Yellow resources have been used by vehicles to transmit BSMs.

As mentioned in Section 2.1.3, LTE-V2X requires a pre-configured, global P value ($P \in \{0, 0.2, 0.4, 0.6, 0.8\}$) to be used by all vehicles as the probability with which they will select

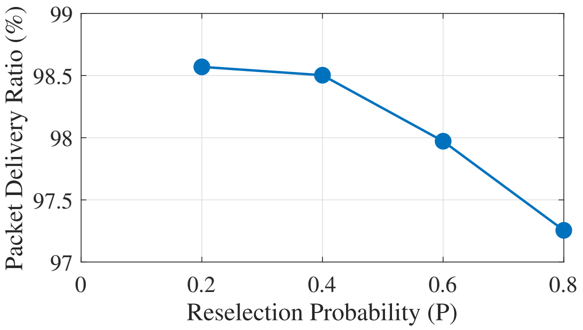

Figure 10: Average PDR for different values of P when no attacker is present.

new resources during reselection (as opposed to retaining their current resources) [19]. We wanted to select P for our simulations such that packet loss would be minimized when no attacker was present (for the reasons expressed above), so we ran several simulations with no attacker while setting the value of P to each of the possible options[8]. These simulations involved 60 vehicles transmitting BSMs for a period of 100 simulated seconds. We calculated the average PDR for each value of P therefrom; as shown in Figure 10, we found that overall PDR decreases slightly for higher values of P. This is logical, since vehicles that reselect resources more often are also more probable to sometimes select conflicting resources. Based on these results, we chose to use $P = 0.2$ in our simulations so as to maximize PDR. Concomitantly, these results allowed us to verify that our LTE-V2X model compares favorably with other simulations of SPS in LTE-V2X from related work (e.g., in [57, 68]), suggesting that our simulation parameters are reasonable with respect to the current literature.

5.5 Detection Through Packet Delivery Ratio

In Section 3, we established that the most common approach to detecting jamming attacks in V2V is through monitoring overall PDR (in this case, of the LTE-V2X channel). We now show, using our LTE-V2X channel model, that PDR-based approaches are unreliable for detecting our targeted sidelink jamming attack because of its minimal impact on overall channel PDR.

[8] For $P = 0$, vehicles never reselect resources, and collisions are extremely numerous since vehicles will never attempt to correct collisions which exist from the beginning. As setting $P = 0$ is unlikely in practice for exactly this reason, we do not consider it here.

5.5.1 Monitoring overall PDR

We consider a hypothetical system monitor who attempts to detect jamming attacks by monitoring the overall PDR of an LTE-V2X channel and raising an alert when that PDR falls below a certain threshold. We assume the monitor can predict the number of vehicles using the channel, and hence the number of BSMs that it should receive in a given time period, with reasonable accuracy. This may, for example, be accomplished by reviewing historical traffic data (e.g., [69]) for metrics like vehicle density, speed, and direction of travel at various times of day and extrapolating therefrom. Alternately, the monitor may make these predictions based on its own long-term monitoring of the channel, e.g., through one or another type of moving average. Based on the number of BSMs it expects to receive for an accurately estimated number of vehicles, we assume the monitor can devise a statistical test and alert to a possible attack if this test is not met by data collected during channel monitoring.

Irrespective of the specific test used to define a PDR-based detection threshold PDR_{TH}, it will always be a function of the number of vehicles n_v using the channel and can be expressed as $PDR_{TH}\left(n_v\right)$. Now, over an interval of t seconds, the monitor will calculate its observed PDR as a test statistic based on n_v, the rate r at which vehicles transmit BSMs, and the number of BSMs b that the monitor receives in that interval. As r and t are known *a priori*, this makes $PDR_{monitor}$ a function of n_v and b as expressed by:

$$PDR_{monitor}\left(n_v, b\right) = \frac{b}{n_v r t} \tag{4}$$

Now, we can say the attack is detectable whenever $PDR_{monitor}\left(n_v, b\right) < PDR_{TH}\left(n_v\right)$; in combination with (4), this yields:

$$\frac{b}{n_v r t} < PDR_{TH}\left(n_v\right) \tag{5}$$

Then, (5), defines an inequality for detection that is useful irrespective of the specific statistical mechanism used to define $PDR_{TH}\left(n_v\right)$. We can take this one step further by defining b more clearly. In our model, packet loss can only result from either packet collisions (due to SPS) or from

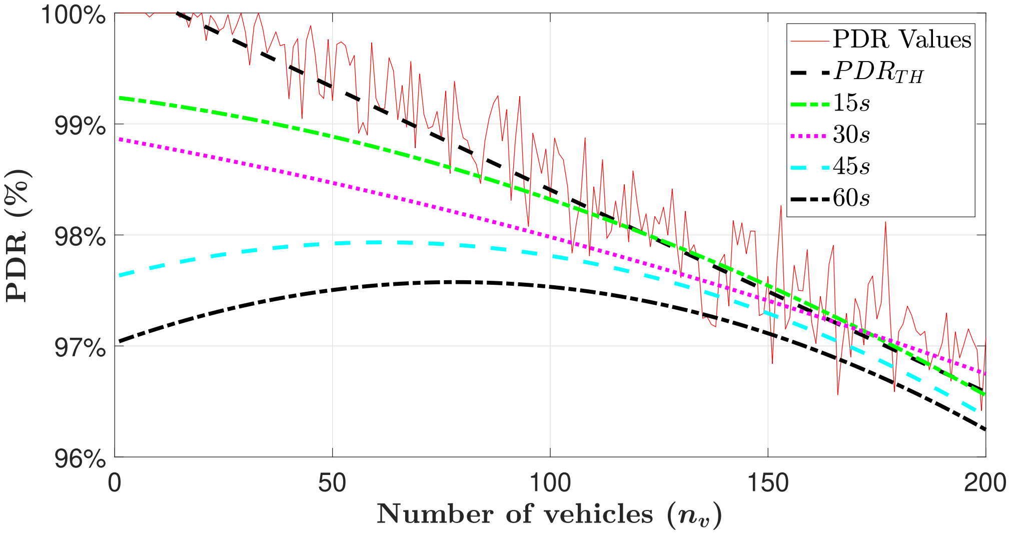

Figure 11: PDR under attacks of varying duration against $PDR_{TH}\left(n_v\right)$.

a DoS attack—in the case of our attack, from a jammer's interference—so we can express b as:

$$b = b_{sent} - b_{SPS}^{lost} - b_{jammed}^{lost} \tag{6}$$

Finally, combining (5) and (6) yields:

$$\frac{b_{sent} - b_{SPS}^{lost} - b_{jammed}^{lost}}{n_v r t} < PDR_{TH}\left(n_v\right) \tag{7}$$

In (7), a relationship is illustrated which is responsible for the unreliability of using PDR as a DoS-detection metric. Note that if b_{jammed}^{lost} were removed from the left-hand side of (7), then the inequality would *always* be true for any reasonable statistical test used to define $PDR_{TH}\left(n_v\right)$. Thus, whether or not the attack is detectable relies solely on the number of messages an attacker jams in an interval of t seconds, leaving it up to the attacker to modify her behavior in order to beat this attempt at detection. Further, (7) illustrates the fundamental problem with using PDR as a detection metric in LTE-V2X systems. Because SPS always causes packet loss, particularly when higher numbers of vehicles are using the channel, the definition of $PDR_{TH}\left(n_v\right)$ must always allow for some level of packet loss (e.g., using standard error or a confidence interval) without raising an alarm; therefore, if b_{jammed}^{lost} is less than or approximate to this necessary allowance for packet loss, the attack will be difficult or impossible to detect.

We can illustrate this problem with an example using our MATLAB model of an LTE-V2X channel. We consider a monitor who defines $PDR_{TH}\left(n_v\right)$ based on mean PDR over time. This monitor tracks mean PDR for each $n_v \in \{0, 1, 2\ \ldots\ , 200\}$ based on observation, calculating a 95% confidence interval on the mean PDR for each n_v. Then, the monitor defines $PDR_{TH}\left(n_v\right)$ as the least-squares regression line for the lower-bound values of the confidence intervals. We ran 1-minute simulations of the LTE-V2X channel using our MATLAB model, calculating the mean PDR and confidence intervals accordingly. This yielded a detection threshold of:

$$PDR_{TH}\left(n_v\right) = -0.0002v + 1.0027 \tag{8}$$

From (8), one can see that the negative slope $\delta = -0.0002$ confirms the expected decrease in $PDR_{TH}\left(n_v\right)$ as n_v increases. This, in turn, means that a DoS attack like targeted sidelink jamming will be less detectable (based on PDR) at higher values of n_v. To demonstrate this, we ran additional 1-minute simulations, this time adding in Eve, who executes the targeted sidelink jamming attack with an (experimentally validated) ability to knock out $14/15$ of BSMs from one vehicle. From (7), detectability is dependent on the number of BSMs jammed by Eve, and hence on the duration of Eve's attack. We therefore evaluated different attack durations of 15, 30 , 45, and 60 seconds and compared the packet loss caused by the attack against $PDR_{TH}\left(n_v\right)$ for all $n_v \in \{0, 1, 2,\ \ldots\ , 200\}$ to see if the effects of the attack can be distinguished from the expected packet loss due to SPS. As illustrated by Figure 11, this is not the case. Particularly for higher values of n_v, PDR during the attack is similar to—and sometimes indistinguishable from—PDR in the absence of an attacker. Note also that even in the worst case we evaluated (a $60\,s$ attack with $n_v < 5$), PDR remains within 3% of expected levels. For all of the attack durations, PDR is generally within 1% of expected levels for all $n_v > 100$. These results illustrate the difficulty of detecting targeted sidelink jamming using PDR—even in our perfect model, the effects of the attack are often indistinguishable from expected levels of packet loss. Therefore, we contend that detection in a more realistic environment (with packet loss from noise, interference, etc.), the attack would be extremely difficult to detect using PDR as a metric.

5.5.2 Monitoring PDR for individual vehicles

An ideal method of detecting our attack would be monitoring the PDR for individual vehicles rather than for the system as a whole. Hypothetically, if a monitor could somehow accurately determine an expected value (and related threshold) for per-vehicle PDR, then it could raise an alert to an attack if the PDR for a particular vehicle fell below the expected threshold. Against this approach, Eve could not hide the impact of her jamming amongst the normal, system-wide packet loss due to SPS like she can to evade detection based on overall PDR. Since the expected PDR for individual vehicles must surely be well above the $7-20\%$ that can be achieved by Alice during the attack - otherwise, LTE-V2X would be useless - such a hypothetical system monitor could surely detect Eve's attack when Alice's PDR suddenly dropped to less than 20%.

Unfortunately, this monitoring approach is generally infeasible in V2V. This is primarily due to the highly unpredictable nature of the V2V environment in combination with the lack of context (to the monitor) for a particular vehicle's movements. A monitor may fail to receive certain BSMs that are blocked from reaching it; for example, what would happen if a large commercial truck pulled over right next to a roadside monitor, creating a large amount of attenuation for signals passing through it to reach the monitor device? Such an incident could easily lead to false positives as the monitor measures a lower PDR for the vehicles whose BSMs are being stopped by this obstacle, does not understand the benign reason for this, and subsequently raises false alarms about an attack. Separately, but similarly, either a static or mobile monitor would suffer from a lack of knowledge about the channel conditions away from its location, and so it might inaccurately expect a low (or high) PDR for *all* vehicles within the entire 1 km communication range it monitors, based only on localized measurements from vehicles that occupy a small fraction of that area.

A secondary barrier to this approach is the lack of identifying information for BSMs at the lower communication layers. LTE-V2X transmissions give no identifying information about the transmitting unit in the lower layers (e.g., there is no MAC address used in LTE sidelink), so a monitor that functions only at the lower layers—which is often desired, if not required, to reduce latency in monitoring systems—will have no ability to determine which BSMs come from which vehicles. To a certain extent, this might be overcome through similar methods as we propose the attacker could use to identify her target in Section 5.1, but it is much less likely that these techniques could feasibly be used to accurately identify the dozens to hundreds of vehicles that could be in range of the monitor at any one time. Also, because vehicles frequently change the resources that they are using, there is no way to establish a probabilistic identification of a vehicle based on its prolonged use of a particular resource. If the monitor were to have access to the upper layer protocols as well, then some identifying information could possibly be used; however, due to the lack of persistent identifiers in V2V (i.e., the exclusive use of pseudonymous identification [38]), even this level of access would be insufficient to *accurately* calculate statistics like PDR for individual vehicles. Therefore, although Eve might have difficulty avoiding detection based on per-vehicle PDR monitoring, we contend that such an approach is unrealistic and we do not concern ourselves further with attempting to avoid this detection technique.

5.6 Better Detection Through Cluster Analysis

Having shown the unreliability and infeasibility, respectively, of attempting to detect targeted sidelink jamming by monitoring overall or per-vehicle PDR, we propose a superior alternative. To this end, it is useful to note that although a vehicle will periodically reselect the resources it uses for transmitting BSMs, there is tendency for the vehicle to select its new resources within the same "frame index" ($SFN\, mod\, 10$) that it is currently using. This occurs in part because new resources are not always reselected (only with probability P) and also because vehicles try to maintain their BSM periodicity across the reselection as close to $100\, ms$ as possible. So, a vehicle is likely to make use of the same SFN index for its transmissions over time irrespective of resource reselections (i.e., a vehicle that enters resource reselection transmitting in the seventh of every ten frames is likely to continue doing so after resource reselection, although it will likely choose different subframes and subchannels *within* that frame). In turn, this means that our targeted jamming attack is likely to impact the same frame index over a significant period of time, leading to frames with that index having lower overall PDR than others. Therefore, a useful detection approach is to monitor PDR changes in specific frame indexes over a sliding time window of ~ 10 seconds and attempt to detect frame indexes with anomalously low PDR values, which may indicate an attack is underway against vehicle(s) that that are using resources in those frames.

Essentially, this approach requires monitoring the channel for 10-second periods and recording how many messages are received in each frame, as well as how many collisions are detected. Related work has shown that detecting packet collisions in V2V is possible [46], and doing so allows the monitor to extrapolate how many messages were actually sent versus how many were received, leading to a PDR calculation for that frame. Then, the monitor can align all of the per-frame PDR measurements by frame index and compare their PDR values to check for anomalies. If a targeted sidelink jamming attack is occurring, then a specific frame index (wherein the victim is transmitting BSMs) should have a much larger number of frames with lower PDR compared to other frame indexes. This is an anomaly detection problem, so cluster analysis is an appropriate and useful technique to apply. Specifically, we propose using the DBSCAN [70] clustering algorithm for this task. DBSCAN is a classic cluster analysis algorithm based primarily on measuring geometric (or related) distance between data points. Compared with other clustering algorithms (e.g., k-Means, hierarchical), DBSCAN is ideally suited for detecting our attack because it can detect clusters of abnormal shape and is specifically designed to identify outliers (i.e., in our case, to identify abnormal PDR levels in certain frames) [71]. DBSCAN is also an unsupervised algorithm and requires neither *a priori* specification of the number of clusters (a limitation of k-Means clustering) nor manual labeling of data, making it a good candidate for autonomous security services to be run on vehicles or standalone roadside infrastructure units. In our case, we propose to use DBSCAN to cluster the PDR values recorded for each frame index and identify outliers; then, if a particular frame index is determined to have significantly more outliers than others, an alert may be raised to a possible attack.

In addition to a data set to analyze, DBSCAN requires three parameters: the minimum number of points to form a cluster (or the "density threshold" [70]), a radius value ϵ to determine which points fall close enough to others to form a cluster, and specification of a function to calculate distance between points. We used Euclidean distance due to our relatively uncomplicated data set (containing only PDR values and their associated frame indexes) and we chose 50 as the minimum number of points to form a cluster based on the dimensionality of our data. As shown in Figures 12a and 12b, we chose an ϵ-value of 0.1 based on sorted k-Nearest distance analysis of our data set,

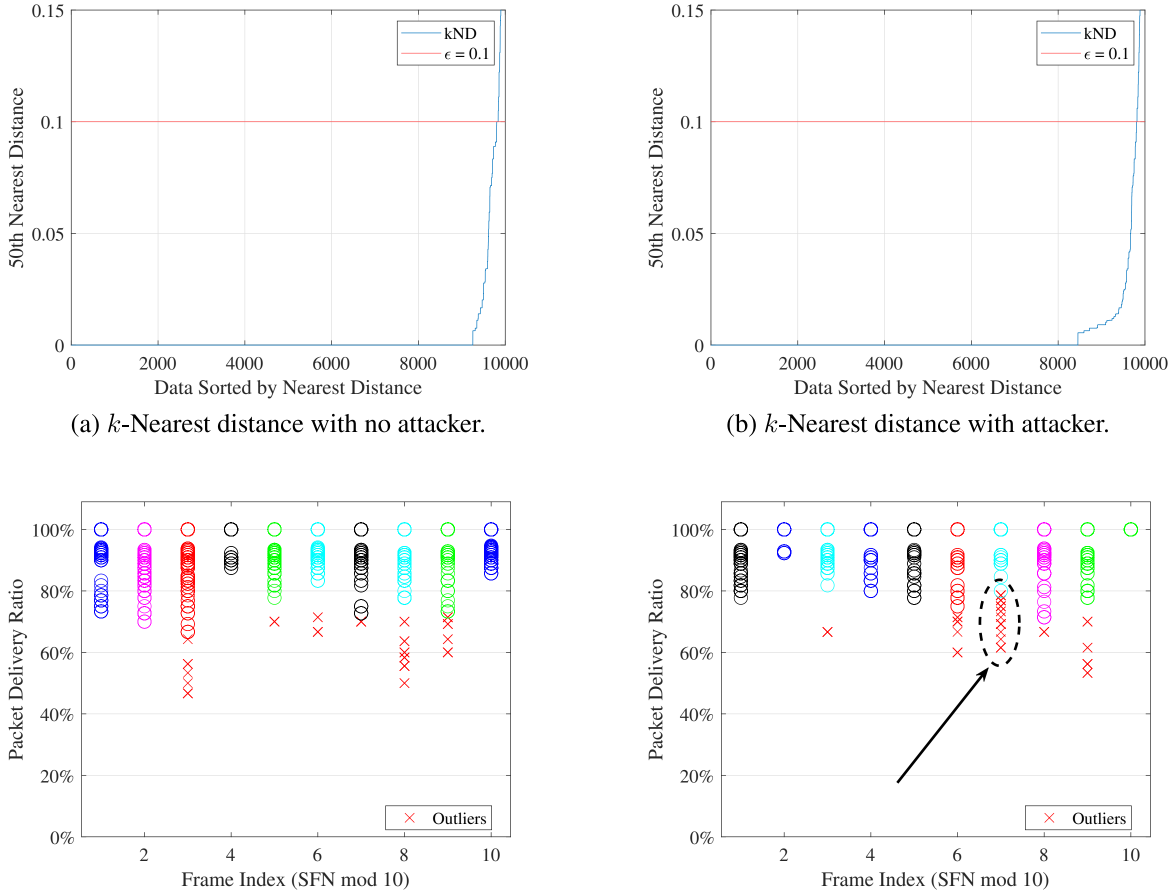

(a) k-Nearest distance with no attacker.

(b) k-Nearest distance with attacker.

(c) DBSCAN with no attacker.

(d) DBSCAN with attacker.

Figure 12: Results showing how DBSCAN can be used to effectively detect targeted sidelink jamming.

which was obtained from simulation of PDR over a 100-second period with 100 simulated vehicles using the LTE-V2X channel. Note that this was a number of vehicles for which we showed PDR monitoring to be unreliable for detecting an attack in Section 5.5. During the simulation, the attacker targeted a vehicle who transmitted on varying resources in every seventh frame. Figure 12d calls out successful identification of the attack. The greatest number of anomalous data points - i.e., PDR measurements lower than normal - are associated with the seventh of every ten frames, in which the attacker was causing packet collisions with its target vehicle. The distinction is visually obvious (another advantage of DBSCAN) when compared with the results in Figure 12c, recorded from the same simulation parameters *sans* attacker. Thus, as desired, we have successfully demonstrated that DBSCAN cluster analysis is a useful approach for detecting our targeted sidelink jamming attack; further, we have shown that DBSCAN detection works where the commonly-used overall PDR metric is highly unreliable. This further addresses our second research question by demonstrating a method of detecting Eve's targeted sidelink jamming attack despite her efforts to blend in with normal system PDR degradation due to SPS.

5.7 Mitigation Techniques

To address the threat we have identified, we propose two mitigation techniques. Each technique has some grounding in related work, and while we leave a complete investigation of their consequences and effectiveness to future work, we contend that these approaches are reasonable and promising in relation to the current literature.

5.7.1 Adjusting the arrangement of sidelink channels

The targeted sidelink jamming attack works because of the precise placement of SCI messages within each LTE-V2X transmission. As described in Section 2.1.2, the control-channel SCI message is always placed in the first two resource blocks of the base subchannel of a transmission; for example, a BSM that is sent using subchannels $3-4$ of an arbitrary subframe will have its associated SCI message placed in the first two resource blocks of subchannel 3. This precise (and predictable) location of the SCI message makes it trivially easy for an attacker to specifically jam the victim's SCI message once that victim's transmission resources (i.e., frame, subframe, and subchannel(s)) have been identified. If this were changed such that the SCI message was not so predictably placed within a subframe, then our attack would be impossible to execute in the efficient, stealthy manner heretofore described. Eve would no longer be able to predict and jam Alice's SCI message, so despite her continuing ability to predict Alice's next BSM, Eve's only chance to block that BSM would be to use a much wider-band jamming signal to target Eve's entire BSM. This would be far less efficient and far more detectable (due to the greater energy required), effectively ruling out this attack for an intelligent adversary like Eve.

The costs of this mitigation technique lie in its effects on the efficiency of receiving LTE-V2X transmissions. Under the current physical-layer procedures [19], receivers can quickly determine which subchannel(s) contain transmissions after receiving a subframe's worth of samples by attempting to decode each location where an SCI message could be located (i.e., at the base of each subchannel) and only proceeding if an SCI message is found. This is an efficient process that allows a receiver to rapidly determine whether there is data to decode in a given subframe and quickly moving on if there is not. Allowing the SCI message to be placed anywhere within a subchannel would require the receiver to attempt decoding every pair of resource blocks in the channel as an SCI message, which would clearly require unacceptably high processing times. The amount of change required to disrupt the attack, though, is not so radical. Merely creating a second allowable placement for SCI messages within a subchannel (e.g., permitting the SCI to be in the first *or last* two resource blocks) would require Eve to either jam both locations, making her more detectable when one of the two jamming messages does not collide with the targeted message and become unrecoverable, or guess which location the SCI is in, significantly reducing her accuracy as jamming BSMs becomes a series of $\frac{1}{2}$-probability guesses instead of the current certainty of success. This would require receivers to check twice as many locations for SCI messages in each subframe, which may still be deemed too expensive in terms of latency, although further investigation (which we leave to future work) would be necessary in order to determine this. In practice, we note that more than two options may be needed; alternatively, this technique may be combined with other proposals to restructure the LTE-V2X control channel (e.g., [56]) by spreading the SCI message across more than one location within the frame. We note also that similar proposals have been credibly made in related work; for example, the performance enhancement for SPS proposed by Mughal *et al.* [56] has a similar potential cost which they argue is offset by the decrease in packet error rate under their scheme. Similarly, the mitigation technique we propose here will result in some additional latency, but we contend that this might be acceptable in exchange for the assurance that targeted jamming attacks like the one we devised cannot be effectively executed without a significantly increased risk of detection. We leave a full investigation of the feasibility and specific implementation requirements of this technique to future work.

5.7.2 Reducing the periodicity of V2V messages

Our targeted sidelink jamming attack takes advantage of the periodicity of V2V messages to predict the timing of several future messages based on receiving just one from an intended target. Therefore, one mitigation approach would be to introduce a level of variability into the precisely periodic nature of V2V transmissions. Under such a system, instead of sending its BSMs exactly every $100\,ms$, a vehicle might slightly adjust its periodicity with each resource reselection. For example, it might send a few BSMs with a periodicity of $98\,ms$, then reselect resources and send its next several BSMs with a periodicity of $103\,ms$. This approach would significantly reduce Eve's ability to predict Alice's messages, as Eve would need to listen and receive two BSMs from Alice before being able to predict subsequent BSM arrivals. If this technique was combined with reducing the intervals between resource reselections (e.g., from every $5-15$ to every $3-7$ transmissions), then Eve's effectiveness would be cut to nearly nothing. In terms of end-to-end BSM latency, this approach would not add any significant delay in getting BSMs from one vehicle to others. However, there are some potentially significant costs to adopting this mitigation technique. For one, related work (e.g. [72]) has shown that reducing the resource reselection interval tends to lead to lower PDR as vehicles become more likely to select conflicting resources when selection occurs more frequently. Second, the impact of varying message periodicity on SPS performance is, to the best of our knowledge, currently unexplored in the literature. However, variable periodicity has been proposed to be used in NR-V2X [8, 72], which uses the same basic SPS algorithm as LTE-V2X, so this may not be a significant concern. Further investigation into the full cost of adopting this mitigation technique is left to future work.

6 Sidelink Resource Exhaustion Attack

In the targeted sidelink jamming attack, the attacker has a very specific target—a single victim vehicle—and her goal is to deny access to the collision-avoidance benefits of LTE-V2X communication for that specific vehicle. We now consider an entirely different attack with a much broader objective. Instead of directly jamming BSMs to prevent them from being received, an attacker might attempt to achieve a similar effect against multiple vehicles simply by making strategic, *prima facie* legitimate transmissions in the LTE-V2X channel. To this end, we propose a *sidelink resource exhaustion* attack wherein an attacker exploits a vulnerability in the SPS algorithm to increase the rate at which other vehicles select conflicting resources and inadvertently "jam" each others' messages. Put another way, Eve attempts to increase the rate of "natural" packet collisions that occurs as a result of SPS by biasing other vehicles' perceptions of channel busyness, tricking them into selecting resources that are likely to conflict with each other, and thereby drastically reducing channel throughput. As we will show, Eve can accomplish this by making only protocol-compliant transmissions, with selectively chosen size and periodicity, in such a manner as to bias other vehicles' choices during SPS resource reselection.

6.1 A Vulnerability in SPS

When a vehicle goes through the resource reselection process outlined in Section 2.1.3, its objective is to select new new radio resources SFN, SLI, and subchannel(s) to use for transmitting its BSMs. An important observation here is that the SPS listening period is $1000\,ms$ (i.e., 1000 subframes, or 100 frames), during which a vehicle builds its set of candidate radio resources CRR. However, the size of CRR is only 10 frames, or $100\,ms$, as this is the defined periodicity for BSMs. Importantly, note that this means the SPS listening window is an order of magnitude larger than the size of CRR in the time domain, meaning that *the value of any particular radio resource in CRR is not based on one, but on many different radio resources* observed during the listening period. This lack of perceptual granularity, i.e., the dependence of each candidate resource in CRR on more than one resource in the listening period, is a serious flaw in SPS which can be exploited by an attacker.

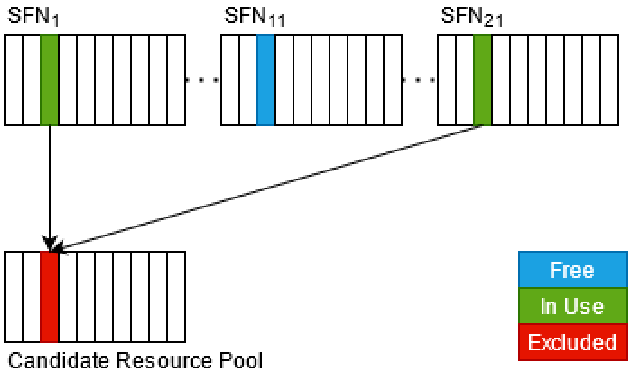

Figure 13: Misalignment in sizes of the SPS sensing window and the candidate resource pool.

In more technical detail, consider that there are ten frames in the listening period for every one in CRR. Thus, as a vehicle processes subframes in the listening period and adds them to CRR, the value that is recorded for that resource's busyness (i.e., whether it is in use or not) is based on an average of ten measurements. An example of this is shown in Figure 13, where the third subframe of the first frame in CRR is marked as in-use, even though it is only sometimes being used. The intent of this design is to cause vehicles to probabilistically choose resources which were observed to be less frequently in-use over the 100 frames preceding resource reselection; however, it operates on the assumption that vehicles are being "honest" and making their BSM transmissions only every $100\,ms$ as expected. Eve's goal is to use this lack of granularity in perception of channel busyness to cause vehicles to misperceive certain resources as being busier than they actually are, thus causing them to improperly exclude those resources as candidates and choose from a smaller candidate resource pool than should actually be considered.

6.2 Attack Idea

During SPS resource reselection (see Section 2.1.3), a vehicle will not select radio resources for which the following conditions are true over the preceding 1000 subframes:

1. A valid SCI message was received in the sidelink control channel.

2. A valid data TB was received in the sidelink shared channel using the subchannel(s) indicated by the SCI message.

3. The average reference signal received power for the TB resources exceeds a defined threshold TH_{rx}.

These requirements (from [19]) preclude, possibly by design, certain attack approaches. For example, Eve cannot simply transmit SCI messages in the control channel and claim to be using resources to transmit data in subchannels that she is not actually using. Other vehicles will know she is not really transmitting data in those subchannels because the received signal reference power for the resource blocks where Eve purports to transmit data will be less than TH_{rx}. Thus, other vehicles will consider claiming those resources anyway, thwarting this naive design for an attack.

In any viable attack, Eve is required to make a complete transmission, using both control and data channels, in any subframe/subchannel resource that she wants other vehicles to perceive as being in use. Note, importantly, that this does *not* mean Eve transmits constantly in the same resources across every frame; our attack is not simply a naive, constant jammer. Also, vehicles are restricted from transmitting more often than once every $20\,ms$ [19]. This requirement is intended to prevent any one vehicle from using too much bandwidth and consequently degrading system availability for all other vehicles (i.e., to prevent exactly the type of effect that we strive to create in this attack). To ensure Eve remains as compliant as possible with LTE-V2X requirements - an important consideration for maintaining her stealthiness - we have made this a restriction of our general threat model as well (see Section 4). As we will show, our attack is able to achieve the effects this requirement was intended to mitigate without violating it or any other LTE-V2X requirements, emphasizing the need for a better PHY/MAC layer design.

6.3 Experimental Validation

In order to ensure our attack would be possible to execute in a real-world scenario, we needed to confirm that real LTE-V2X equipment would respond as expected to Eve's transmissions. Specifically, we wanted to confirm that periodic, standards-compliant transmissions by Eve would be able to influence the resource reselection procedure of commercial LTE-V2X equipment. We evaluated this by setting up an experiment with three devices: one Cohda device, transmitting BSMs at the standard $10\,\mathrm{Hz}$ rate with a transmit power of $23\,\mathrm{dBm}$, one USRP B210 as the attacking device, transmitting sidelink frames on configurable time/frequency resources with $10\,\mathrm{dBm}$ transmit power at $20\,ms$ intervals, and another USRP B210 running a sidelink receiver (adapted from srsRAN [62]) as a channel monitor to observe the impact of Eve's transmissions on the Cohda device's use of radio resources.

The objective of this experiment was to show that it is possible for a device like a USRP B210, whose maximum transmit power of approximately $10\,\mathrm{dBm}$ is well below the $23\,\mathrm{dBm}$ maximum set for Eve (and the $23\,\mathrm{dBm}$ Cohda transmission power), to influence the resource reselection process of the Cohda device. This goal had to be shown indirectly because the Cohda devices use a proprietary, kernel-level functionality for SPS; therefore, there is no way to directly view their perceptions of channel usage during reselection. To evaluate Eve's impact, her USRP B210 was configured to transmit one BSM on subchannels $1{-}2$ of the first subframe in every other frame (i.e., every $20\,ms$). Our experimental objective was to determine whether the Cohda device would avoid using subchannels $1{-}2$ of subframe 1, not only in the frame that Eve transmits in but in *all* frames. Note that this means avoiding not only the subchannel pair $1{-}2$ but also the pair $2{-}3$, as the latter selection by the Cohda device would also conflict with Eve's use of subchannels $1{-}2$.

Due to the inherent randomness of SPS resource reselection, we needed to make sure we ran this experiment for a sufficiently long period of time as to be certain that if the Cohda device did not select the resources used by Eve, this was due to the effect of Eve's attack rather than random chance. Given a $10\,\mathrm{MHz}$ channel with 5 subchannels, there are 4 pairs of adjacent subchannels $(1{-}2, 2{-}3, 3{-}4, 4{-}5)$ from which one can be selected to carry a 2-subchannel BSM. With 10 LTE frames in a $100\,ms$ BSM interval and 10 subframes per frame, this gives $10 * 10 * 4 = 400$ total candidate resources to select from. The chance of the Cohda device selecting any particular resource (i.e., any pair of adjacent subchannels) is therefore $\frac{1}{400}$, and the chance of selecting either of the two resources that Eve's transmissions cover is

$$P(R1 \vee R2) = \frac{1}{400} + \frac{1}{400} = \frac{2}{400} = 0.005$$

where $R1$ and $R2$ are subchannel pairs $1{-}2$ and $2{-}3$, respectively. Thus, the chances of a device *not* selecting either of those resources is $\neg P(R1 \vee R2) = 0.995$. Recall that for each reselection, a vehicle may keep its existing resources and only reselect with probability $P(reselect)$, so the chance of a vehicle choosing those resources for a given reselection is actually $P(reselect) * 0.995$. The actual $P(reselect)$ used by the Cohda devices is unknown, so assuming a worst-case value of $P(reselect) = 0.2$, the probability of a device not choosing either of the two resources used by Eve can be formulated as:

$$\begin{aligned}
&\neg P\left(reselect\right) \vee \left[P\left(reselect\right) \wedge \neg P\left(R1 \vee R2\right)\right] \\
&(1 - 0.2) + (0.2 * 0.995) \\
&0.8 + (0.199) = 0.999
\end{aligned}$$

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 that a device performs resource reselection at most once every $1500\,ms$ (i.e., after 15 transmissions), the experiment needed to be run for $4603 * 1500 = 6.9e6\,ms = 115.7$ minutes in order to claim, with 99% confidence, that the Cohda device's avoidance of Eve's resources was due to her transmissions rather than simple probability.

Therefore, the experiment was run for 116 minutes to meet this level of confidence. Results from the monitor device, as shown in Figure 14, confirmed that the Cohda device did indeed avoid choosing the resources used by Eve over a 116-minute period. Although it selected every other resource at about the same rate, the Cohda device did not choose subchannel pair $1{-}2$ or $2{-}3$ of subframe 1 (called out in red in Figure 14) due to Eve's use of those resources in just 50% of frames. This validates the central premise of our attack by demonstrating the ability of an attacker to influence a commercial LTE-V2X device's selection of radio resources with nothing more than standard-compliant, periodic transmissions in the LTE-V2X channel.

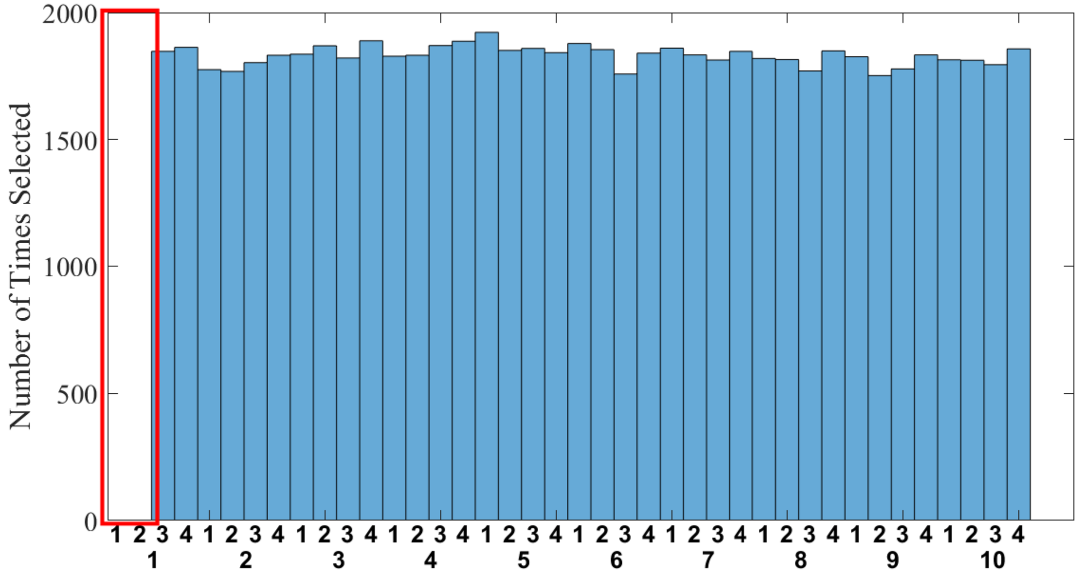

Figure 14: Resources selected by the Cohda device during a 116-minute sidelink resource exhaustion attack. Values on the x-axis are candidate resources (subchannel pairs) grouped by subframe $1-10$.

6.4 Attacking the LTE-V2X Channel

Although our hardware experimentation was necessarily limited to an indirect proof of our premise, we can more thoroughly evaluate the effectiveness of our attack against an LTE-V2X channel using MATLAB simulation. Our objective here is to extrapolate from our experimental results to show that Eve can have the same impact on multiple vehicles as we showed on one, to the extent that vehicles begin avoiding so many resources that they begin selecting conflicting resources and interfering with each others BSMs. To accomplish this, we use the same MATLAB simulation environment described in Section 5.4. We ran simulations for different approaches that Eve might take; for example, we sought to determine how Eve's use of varying transmission size (i.e., number of subchannels) and transmission periodicity would improve or decrease her success. We also varied the number of simulated vehicles using the channel to determine whether Eve's attack is more effective when the channel is under certain levels of use.

While we evaluated how changes in Eve's transmission periodicity and size impacted her effectiveness, we configured her to reselect resources almost as often as she is allowed to. In a real LTE-V2X environment, Eve's continued use of the same subframe and subchannel over time would

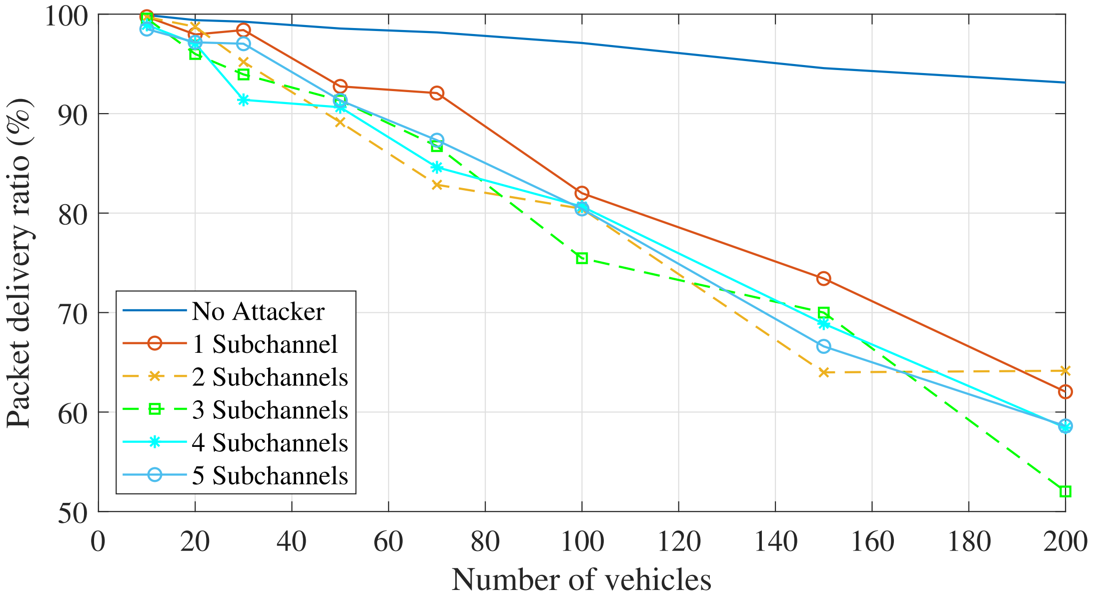

Figure 15: Effects of the attacker's transmission size on effectiveness of sidelink resource exhaustion.

cause direct packet collisions, which is not her intended approach, so we allow (and require) her to change resources in such a manner as to avoid packet collisions whenever possible. Technically, it is at each vehicle's discretion (e.g., for latency requirements) to reselect resources at any rate up to and including after every transmission [19]; however, we configured Eve to change resources after every $2{-}3$ transmissions instead as we consider changing resources after every transmission to be unrealistic. We thus ensured that Eve was configured to avoid using the same resources as were being used by any other simulated vehicle, affirming that any observable impacts on the channel result from the attack rather than from Eve inadvertently jamming vehicles' messages directly.

In our first round of simulations, we held Eve's transmission periodicity constant at $20\,ms$ and varied her transmission size between $1{-}5$ subchannels. Eve's goal is to prevent other vehicles from using as many resources as possible, not only in the frames she transmits in but in all frames (thus exploiting the vulnerability in SPS described in Section 6.1). The results for these simulations are shown in Figure 15. Compared with the control simulation (i.e., the simulation with no attacker), which is shown as a solid blue line in Figure 15, we found that Eve's attack reduced overall PDR irrespective of the number of vehicles using the channel, although the most significant impacts occurred when a greater number of vehicles were present. This makes sense, as Eve's use of more subchannels for her transmissions means other vehicles believe there are fewer unused resources

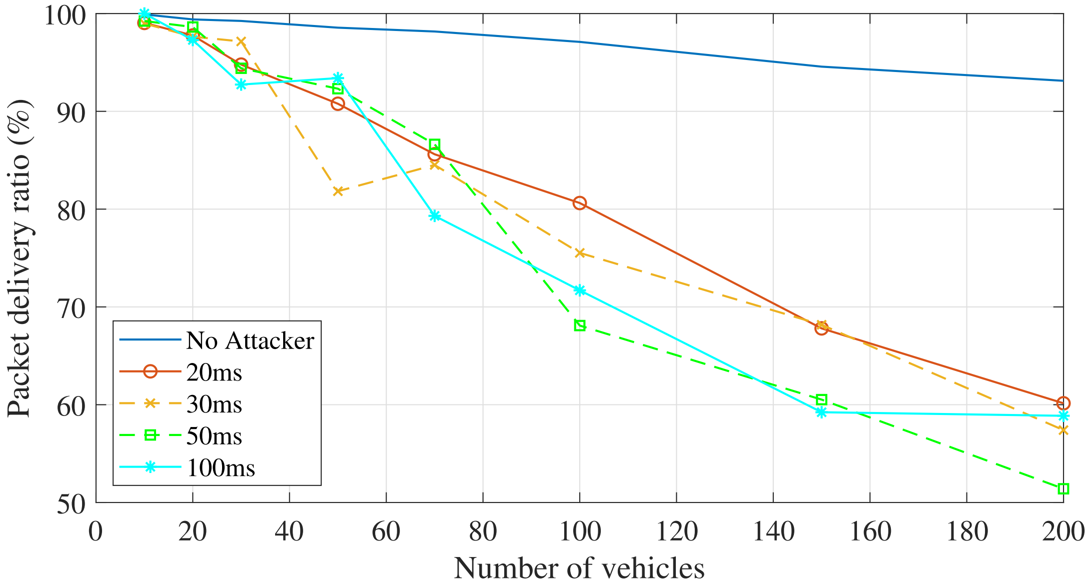

Figure 16: Effects of the attacker's transmission periodicity on effectiveness of sidelink resource exhaustion.

to choose from, causing them to select resources from a smaller pool and consequently to choose conflicting resources with greater frequency. The results in Figure 15 show that just by transmitting normal, 2-subchannel BSMs with $20\,ms$ periodicity, Eve is able to cause an overall PDR reduction of nearly 40%; with slightly larger messages, Eve can have an impact approaching a 50% reduction in overall PDR. This demonstrates both that sidelink resource exhaustion is effective and that its effectiveness stands largely irrespective of Eve's transmission size.

We also evaluated how Eve's transmission periodicity impacts the effectiveness of her attack. This is important to consider because we presume that Eve will want to transmit as infrequently as possible to minimize her chances of being detected. So, we performed additional simulations to evaluate Eve's effectiveness at different transmission periodicities, holding transmission size constant at 2 subchannels to imitate a legitimate vehicle sending standard-size BSMs. Figure 16 shows the results of these simulations. We configured Eve to vary her transmission periodicity between $20, 30, 50$ and $100\,ms$ (the acceptable choices for LTE-V2X BSM periodicity [7, 67]). Our results, shown in Figure 16, show Eve's ability to cause severe degradation in PDR regardless of her transmission periodicity. In fact, for a periodicity of $50\,ms$, Eve is able to reduce PDR by nearly 50%. Figure 16 shows that once again, Eve's effectiveness increases with the number of vehicles using the channel; as before, when more vehicles are present and must select resources from the same perceptually diminished resource pool, they tend to select conflicting resources with greater frequency.

One might express concern that our attack is less effective when the channel is less busy. This is an unavoidable consequence of our attack design, as its primary mechanism is abusing SPS to increase the probability with which multiple other vehicles select conflicting resources. Thus, when fewer vehicles are present, the attack is inherently less effective (although not impotent). Two factors should be considered to contextualize this apparent shortcoming. First, our results come from an idealized channel model where SPS resource selection conflicts are the only causes for packet loss. In a real environment, packet loss for any number of vehicles will be significantly higher than our idealized baseline due to the harsh conditions of a noisy, dynamic V2V channel [57]. 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 environment factors like multipath fading, noise, and others. Second, our results show that Eve's impact is inversely related to the number of vehicles using the channel. When more vehicles are present, SPS is supposed to allow the channel to balance the load effectively so that vehicles more or less end up fully using all channel resources, with each vehicle using no more than one BSM slot per period in the worst-case scenario to minimize BSM latency. However, we have demonstrated that the busier the channel is, the greater havoc Eve can inflict; thus, our attack reduces LTE-V2X channel throughput under exactly the conditions when it has the greatest need for maximal efficiency.

6.5 Efficiency and Effort Considerations

As we did with the targeted sidelink jamming attack, it is useful to discuss the effectiveness and efficiency of this attack in a real-world environment.

6.5.1 Efficiency

Evaluating the duty cycle of our sidelink resource exhaustion attack reveals it to be highly efficient, particularly when greater numbers of vehicles are using the channel. For this attack, because it is against the entire channel rather than just one vehicle, it makes sense to talk about duty cycle in terms of both frequency and time. In a $100\,ms$ BSM period, there are $100 * 50 = 5,000$ resources that can be used. Eve transmits every $20\,ms$ using a variable number of subchannels, but we will use 5 here as a worst-case argument for calculating the maximum duty cycle. Thus, Eve transmits in $5 * 20 = 100$ resources per BSM period. As shown in Figure 15, this achieves a reduction in the number of used channel resources of about 70% in a very short period of time. Therefore, Eve's use of 100 resources per BSM interval results in an effective elimination of $3,500$ of $5,000$ resources in subsequent intervals, a duty cycle of just 2.9%. This is highly efficient attack which, as Figure 15 illustrates, only gets more efficient as time elapses.

6.5.2 Attack range

Another important consideration is the effective range of Eve's attack. As we have shown, sidelink resource exhaustion is far more effective when larger number of vehicles are using the LTE-V2X channel, meaning it is necessary that a sufficiently large number of vehicles are within range of Eve's transmissions for the attack to be as effective as desired. Since Eve is behaving in the same manner as other vehicles (excepting her resource reselection interval and non-random choices during resource reselection), we can reasonably say that her transmissions impact other vehicles' SPS decisions within the same range that genuine transmissions would—of course, this is the key insight behind our attack. One element of the SPS algorithm not heretofore discussed is the third criterion, which states that a resource shall be excluded from CRR if, "[t]he average reference signal received power for the...resources exceeds a defined threshold TH_{rx}" [19]. From this, we can conclude that any vehicles affected by Eve must be within a distance where TH_{rx} is exceeded by Eve's signals when they arrive. Unfortunately, the standard does not specify a value for TH_{rx} and leaves it to be defined on an implementation-specific basis. Based on LTE-V2X field testing [73], a $-95\,$dBm received power level is sufficient to successfully recover an LTE-V2X transmission, from which we can use (1) and a link-budget calculation to determine that Eve should be able to impact any vehicle within approximately $8.06\,$km. Now, as this is roughly 8 times the expected communication range of LTE-V2X, it is important to point out that when factors beyond free-space path loss (e.g., multipath fading, noise, interference) are considered, this distance will be considerably less. However, this calculation is sufficient to underscore the main point, which is that Eve should be able to impact any vehicles that are within the normal LTE-V2X range of $1\,\text{km}$. A more specific description of Eve's range would need to be performed on an implementation-specific basis where the value of TH_{RX} is known, which would facilitate a precise calculation of Eve's effective range and any required adjustments to compensate for TH_{RX} increasing during the execution of SPS. Such a further investigation is left to future work.

6.6 Detection Strategies

We can now turn to answering our second research question with respect to the sidelink resource exhaustion attack, which we do by discussing how detectable Eve is during execution of the attack. As we did with our first attack, we begin by examining how and why the state-of-the-art approach to BSM jamming detection, which is based on monitoring overall PDR, is not an effective approach to detecting our resource exhaustion attack. We then propose a superior alternative based on least-squares regression analysis and show, through simulation, that it can effectively detect our attack where state-of-the-art techniques cannot.

6.6.1 Naive observation of vehicle transmission patterns

An intuitive approach for detecting our resource exhaustion attack relies on observing when a vehicle's transmission patterns do not align with expectations, which may be based either on operating standards (e.g., [19]) or on a hypothetical system monitor's real-time monitoring of an active system. For example, an observation that periodic transmissions are using up the same n resources of every other frame might be suspicious, especially if this pattern continued over a significant period of time, as such a usage pattern would rarely emerge naturally from fair and honest use of SPS by all vehicles. This monitoring approach may allow easy detection of Eve if she uses certain naive attack parameters, e.g., transmitting larger-than-average messages at short intervals in the same subframes for a long period of time. However, our Eve is an intelligent attacker who does not facilitate such easy detection of her actions. If Eve breaks up her transmission pattern (e.g., by varying transmission size and subframe as frequently as allowed), then she will be more difficult to detect. BSM size is dynamic in reality; thus, Eve's regular changes of message size would not be suspicious. Similarly, since BSMs are allowed to be sent more frequently than the regular 10 Hz rate [67], Eve's transmissions at 20, 30, or $50\,ms$ intervals[9] are equally acceptable.

Eve's compliance with periodicity and message size requirements mean she can execute this attack while remaining *syntactically* compliant with the LTE-V2X standards. However, it is equally important that Eve remain *semantically* compliant with V2V protocols. Her compliance with required power levels, transmission periodicity, etc. is irrelevant if the messages themselves contain nonsense data or obviously falsified information, as it would be easy for receivers and any form of system monitor to pick up on regular transmissions of gibberish as a suspicious goings-on. However, if Eve sends genuine V2V BSMs, she may be required to identify herself in some way, e.g., for authentication purposes. This is a risky proposition because, while doing so may get past initial semantic message checks, if her traffic is ultimately identified as malicious than she will have identified herself to the authorities in the course of trying to avoid detection. There are a variety of ways that Eve may attempt to avoid detection by a semantic system monitor, but most of them fall outside the scope of this thesis; for example, abuse or theft of credentials for an upper-layer security protocol might be a way around such monitoring. The exploration of how upper-layer detection and evasion may play out in detail is left to future work.

6.6.2 Monitoring overall resource usage patterns

A more intelligent system monitor may attempt to detect our resource exhaustion attack by observing overall radio resource usage patterns rather than looking for irregularities in the usage of specific resources. In the frequency domain, a monitor may observe each subchannel over time to verify that they are all being used equally, or at least in consistently unequal ratios[10]. Since Eve's overall

[9] Eve will also not be the only vehicle changing her transmission periodicity. 3GPP TS 36.213 [19] permits—in fact, requires—vehicles to alter their periodicity as needed to meet requirements for latency, error correction, quality of service, etc.

[10] Given the most common BSM size of 2 subchannels, subchannels $2-4$ are naturally going to be more often used than subchannels 1 or 5 because multiple pairs of adjacent subchannels overlap the middle three, but only one pair overlaps each of subchannel 1 $(1-2)$ and 5 $(4-5)$

objective relies on transmitting in specific subchannels, aiming to make them seem more used than they really are, a monitor may observe all subchannels to see if any particular subchannel(s) are being abnormally underutilized. For example, if subchannel 2 is only being used for $2-3$ transmissions per frame over a sustained period of time when other subchannels are being used completely, the monitor might become suspicious that vehicles are not selecting subchannel 2 despite its general availability.

Eve may attempt to circumvent this detection approach by regularly changing the radio resources that she wants to make appear in-use. Since all vehicles in the system reselect resources at most once every $1.5\,s$ (15 BSMs at $100\,ms$ intervals), and also because the listening period for reselection looks back only $1s$, Eve can ensure she influences every vehicle's reselection process by targeting certain resources for exclusion over a $\sim 3\,s$ interval. Thus, all vehicles may select from one diminished resource pool for a few seconds, then from a different, similarly diminished pool, and so on. A monitor might notice brief periods when one or more subchannels are underutilized, but the lack of persistent underutilization of any specific resources would preclude straightforward identification of anomalies that would be indicative of this attack. So, attempting to detect sidelink resource exhaustion by monitoring the usage of specific resources (i.e., subchannels) is unlikely to be a viable technique.

6.7 Detection Through Regression Analysis

Although monitoring the usage of specific subchannels (or subframes) over time is not a workable solution for detecting our attack, monitoring the overall number of resources used over time is likely to be a more fruitful approach. By examining channel resource usage levels over time, it is possible to observe when the pool of used resources appears to be diminishing for unknown reasons—i.e., because a sidelink resource exhaustion attack is underway. This approach is particularly useful over periods of time when the number of vehicles remains relatively constant because when this condition holds true, the general levels of resource usage in the channel should remain consistent. That is to say, if n vehicles are using the channel over a period of one minute, one would expect that while they will use different resources over that period (due to many iterations of resource reselection),

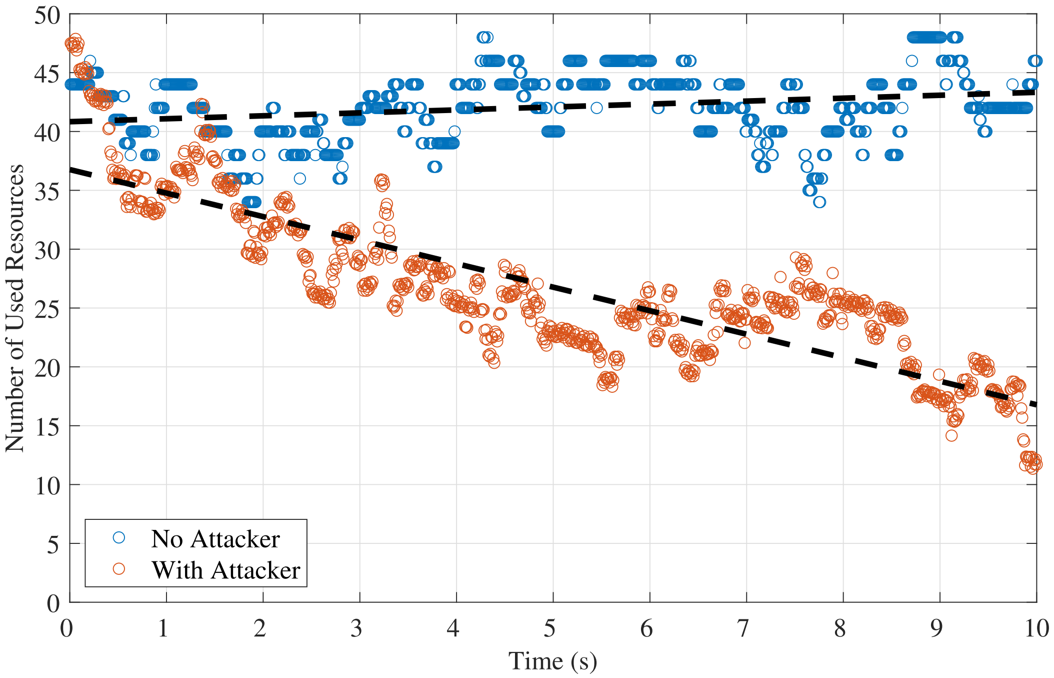

Figure 17: Trends in channel resource usage over time when the channel is under attack versus normal operation. Least-squares regression lines (black and dashed) show clearly divergent trends between these scenarios, facilitating detection of the attack.

the *total number of used resources* should be nearly constant. If, instead, a monitor observes fewer and fewer resources being used as time goes on—something that, by design, never occurs under normal SPS operation—then this would be a strong indicator that a resource exhaustion attack is underway. One way to monitor channel resource usage in this fashion is to approximate trends in channel usage over time using least-squares regression analysis, thus allowing an alert to be raised if overall resource usage trends downward significantly (in a statistical sense) over time.

To evaluate the effectiveness of this approach, we ran the same simulations as before (varying Eve's transmission sizes and periodicity for different numbers of vehicles), but with a simulated monitor set up to record the total number of used resources in each frame over time. Then, we applied least-squares regression analysis to the collected data. In Figure 17, the results are shown for the simulation with no attacker as well as one of the simulations with an attacker present. The results shown in orange were collected during the attack with 100 simulated vehicles using the channel. As indicated by the black, dashed least-squares regression lines on each scattered data series, there is an obvious divergence of trends in resource usage when an attack is or is not ongoing.

In the absence of an attacker, the slope of the least-squares line is $\delta = 2.4e{-}3$; as expected (see above), this is a negligible change. On the other hand, when an attacker is present, regression yields a strong negative trend of $\delta = -0.2$. More in-depth analysis and experiments with real equipment will be required to establish a threshold for when the trend in resource usage is sufficiently negative to deem indicative of an attack; however, it suffices here to note that the clearly divergent regression lines shown in Figure 17 support the validity of this technique for detecting our attack.

As a final point of discussion, one might legitimately wonder how Eve is able to have such a significantly negative impact on resource usage over time, since she *should* only be able to impact the channel usage levels inasmuch as any single vehicle can. Two factors are involved in explaining these results. First, Eve is transmitting more frequently—importantly, though, at an entirely legitimate and unsuspicious rate—than other vehicles generally will. This allows her to make certain subchannels seem busier than they are, causing vehicles undergoing resource reselection to avoid those subchannels. Second, Eve's frequent and deliberate reselection of resources allows her to cause a sort of cascade effect that pushes other vehicles eventually towards using fewer and fewer resources. Figure 18 depicts how this works using a simplified channel model with three subchannels and five subframes per frame. An arbitrary number of vehicles are making transmissions, which are shown in green, while Eve's transmissions are indicated in red. At the top of Figure 18, the normal behavior of SPS is shown. After the listening period, a candidate resource set has been created, with the resources used least frequently (i.e., those from which new resources will be selected) shown with green highlighting. With no attacker, candidate resources are spread across all subchannels. However, once the attack begins, this changes rapidly. As Figure 18 shows, Eve's transmissions at first push vehicles away from using the top subchannel. Then, after several vehicles reselect resources based on this biased candidate resource set, observe that Eve's transmissions, which are now in the middle subchannel, cause the resulting candidate resource set to have low-use resources almost exclusively in the top subchannel. After one additional iteration, all of the candidate resources with low usage are in the same subchannel. By the end of this simplified example, every vehicle that needs to reselect resources will attempt to choose new resources in the same subchannel, illustrating how the used bandwidth narrows over time during the sidelink resource exhaustion attack. Figure 18 is simplified out of necessity—depicting the real process,

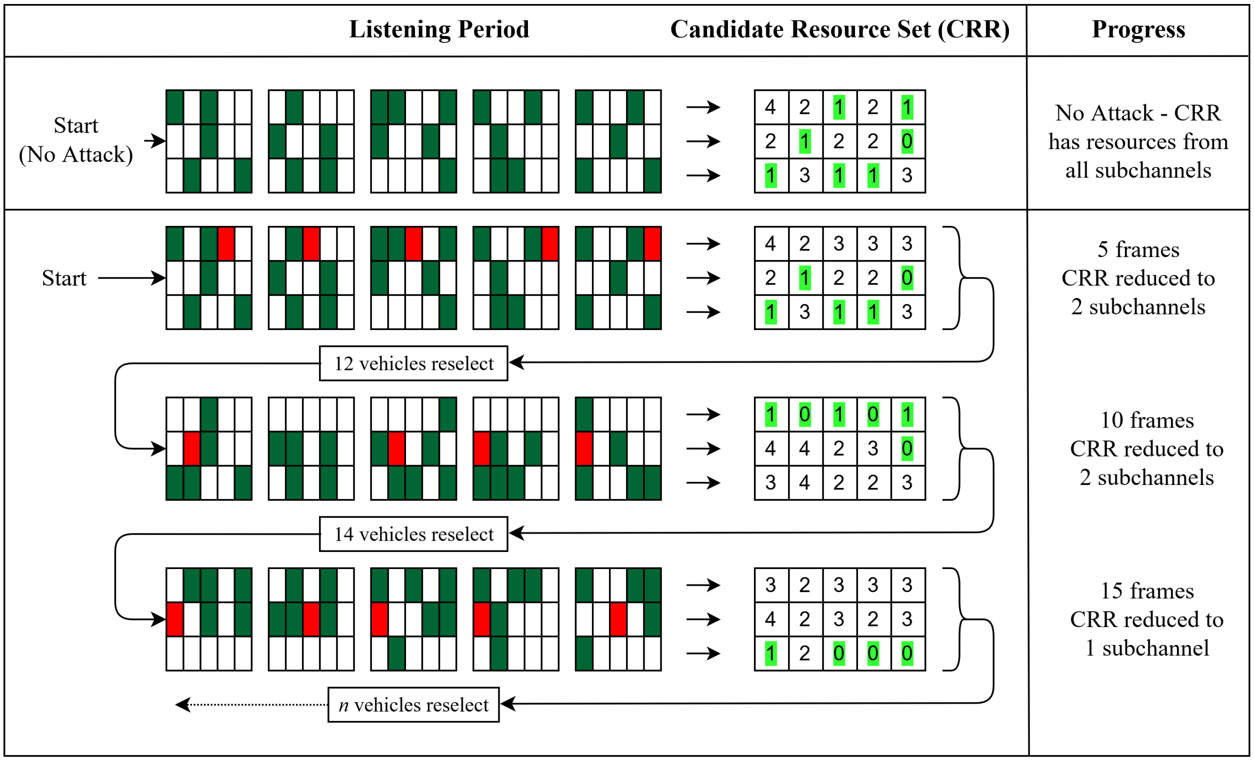

Figure 18: Visualization of the mechanics behind sidelink resource exhaustion.

which occurs over hundreds of frames, will not fit within the confines of a page—but its point applies equally to the full-sized LTE-V2X channel. As the illustrated process continues, Eve is eventually able to exclude not only one, but first two, then three, and possibly even four out of five (real) subchannels from being used by other vehicles. This is the reason for the observable decline in overall channel usage during the attack which can be seen in Figure 18, as well as the core reason why the attack is effective at causing a decrease in PDR, since the reduction in the number of used subchannels in a real system inevitably leads to some vehicles choosing conflicting resources and inadvertently jamming each others' BSMs.

6.8 A Mitigation Proposal

Our sidelink resource exhaustion attack exploits a fundamental vulnerability in SPS, the misalignment in size between the listening period and candidate resource set. One possible mitigation for our attack, then, is addressing this vulnerability by modifying the SPS algorithm so that this misalignment 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 candidate resource set). Doing so would eliminate the possibility of earlier, irrelevant transmissions (including the attacker's) from influencing a vehicle's choice of resources; however, it would also reduce the vehicle's ability to accurately estimate channel busyness. The costs of this have not, to the best of our knowledge, previously been studied. We suppose such costs might include a higher rate of packet collisions due to vehicles being less aware of who is using the channel when selecting new resources, due to the potential that this reduced awareness could increase the rate at which vehicles choose conflicting resources. However, a full investigation of this avenue of approach requires further study and is left to future work.

The alternative approach is expanding the size of the candidate resource set to $1000\,ms$ to match the SPS listening period. This bears a certain similarity to proposals for SPS performance enhancement such as [57, 68, 72]. A common theme in those and other related works is the idea of extending the period between resource reselections to reduce the chance that vehicles choose conflicting resources (a consequence of reducing the number of reselections in general). 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. In doing so, we eliminate the ability of Eve to improperly constrain resource selections, as her rate-limited transmissions will not have nearly such an effect on a larger candidate resource set as they can on the current configuration. One notable consequence (and potential downside) of this approach is that the interval between the last BSM preceding reselection and the first to be sent afterwards would be much greater than the standard $100\,ms$, ranging in fact up to a full $1\,s$ between messages. On average, we expect the interval to be closer to $500\,ms$, but this is still a five-fold increase in latency for one out of every $5-15$ messages. Whether or not this is an acceptable price to pay requires further investigation, particularly concerning real-world traffic flows and whether an intermittently longer delay between BSMs is likely to significantly increase the likelihood of an otherwise avoidable collision occurring. We leave this further investigation to future work.

7 Comparing Our Novel Denial-of-Service Attacks

We have presented two novel DoS attacks against LTE-V2X: targeted sidelink jamming and sidelink resource exhaustion. In presenting each attack, we discussed real-world considerations like jamming-to-signal ratio (JSR), effective range, and power efficiency in Sections 6.5 and 5.3.1. Now, we can consider the respective strengths and weaknesses of each attack compared with the other.

7.1 General Effort

In Section 5.3, we determined that our targeted sidelink jamming attack has a duty cycle of 10%. This is *prima facie* a fairly inefficient attack, especially when compared with other examples from related work. However, one must also consider that this relatively high duty cycle is a trade-off for deniability. Jamming attacks of this type (i.e., which aim to increase the noise level for a specific signal so as to make it irrecoverable by a receiver) with very low duty cycle often require actions which are drastically different from normal transmission (e.g., sending a high-power, narrow-band jamming signal for a very short interval). Our attack, in contrast, has the advantage of being deniable in the context of SPS (as we mentioned in Section 5.1), so detection techniques that alert on abnormal energy levels in the channel will not be effective. That said, the efficiency of our targeted sidelink jamming attack pales in comparison to that of sidelink resource exhaustion. We calculated that the latter has an impressive duty cycle as low as 2.9% (see Section 6.5), making it far more efficient with respect to this metric. Also, sidelink resource exhaustion is able to inflict serious DoS effects on a large number of vehicles, whereas targeted sidelink jamming only affects one targeted victim. Which of these is more desirable will depend heavily on the motivations of each specific attacker; if mass disruptions in the LTE-V2X channel are desired, then sidelink resource exhaustion is the way to go, but targeted sidelink jamming is a far better (and less overt) approach for silencing a specific victim's LTE-V2X transmissions. One might also argue that an attacker who is sufficiently motivated to single out a particular vehicle for an attack might find the extra effort required to be acceptable, although this would be a subjective choice. Overall,

Table 3: Efficiency comparison of targeted sidelink jamming and sidelink resource exhaustion.

	Duty Cycle	P_{TX}	Effective Range
Targeted Sidelink Jamming	10%	$10 - 23\,\text{dBm}$	Varies with P_{TX}
Sidelink Resource Exhaustion	2.9%	$23\,\text{dBm}$	$\sim 1\,\text{km}$

each of our attacks has trade-offs between efficiency and objectives which must be balanced on a case-by-case basis for an attacker to decide their approach.

7.2 Power Efficiency

From the perspectives of required transmit power and length of the attack, targeted sidelink jamming is a more flexible and generally more efficient attack than sidelink resource exhaustion. As we discussed in Section 5.3, targeted sidelink jamming is effective whenever the JSR is at or above $0.04\,\text{dB}$, meaning that the jamming signal will effectively jam BSMs at any receiver where Eve's and Alice's signals arrive with approximately equal power. Achieving this JSR leaves Eve with much discretion over her actions, allowing her to customize her transmit power, attack duration, and distance from Alice as needed for specific situations. For example, Eve might carefully position herself along a road where she knows Alice regularly travels in such a manner as to jam messages over a very small area for the short period when Alice is present (e.g., at a blind intersection that Alice must pass through) using minimal transmit power. Alternatively, Eve could follow a car-length's distance behind or beside Alice and transmit with equal power to her over a longer period of time, effectively jamming a wider area at the price of lower power efficiency. These examples illustrate the range of options that Eve has with respect to power efficiency when executing the targeted sidelink jamming attack.

In contrast, Eve has less discretion with her transmit power level and attack duration when executing a sidelink resource exhaustion attack. We showed in Section 6.4 that the attack is most effective when more vehicles are in range of each other and using the channel, strongly motivating Eve to maximize her transmit power in order to affect the greatest number of victims. Thus, Eve is generally likely to use her maximum allowable transmit power of $23\,\text{dBm}$, with the exception of limited circumstances where vehicle density is extremely high (e.g., in a major city center, on an interstate at rush hour). We also showed in Section 6.4 that the attack becomes effective over time, requiring Eve to continue transmitting in order to maintain the effectiveness of the attack. Comparatively, then, targeted sidelink jamming is likely to more often allow the attacker to accomplish her objectives with lower transmit power levels and a shorter attack duration than sidelink resource exhaustion, giving the upper hand to the former attack with respect to this metric.

7.3 Required Distance for Effectiveness

Sidelink resource exhaustion has a range of effectiveness that depends solely on the transmit power chosen by the attacker, but generally covers the standard LTE-V2X communication range of 1 km. Targeted sidelink jamming has a range of effectiveness that is also dependent on the transmit power chosen by the attacker; however, the effective range is based on JSR, which is also dependent on the distance from the victim. Thus, targeted sidelink jamming has a more limited range of effectiveness and can impact receivers which, generally speaking, are closer to the attacker than the victim (although the specific delineation of the jamming area is dependent on the attacker's transmit power). All of that said, though, comparing the effective distance for each attack is of limited usefulness due to their divergent objectives. Sidelink resource exhaustion is more effective when more vehicles are in range and is intended to affect all channel users, so a larger range of effectiveness is both necessary and desirable. In contrast, targeted sidelink jamming is directed against a single victim, in response to whose movements the attacker can customize power and positioning decisions, so a larger range of effectiveness is neither necessary nor (potentially) desirable. In summary, sidelink resource exhaustion is superior in the strictest sense that it has a broader range of effectiveness; however, the usefulness of this conclusion is limited by the context of attacker intent.

8 Conclusions and Future Work

In this thesis, we identified fundamental vulnerabilities in the PHY and MAC layers of LTE-V2X and devised novel, intelligent DoS attacks to exploit them. We set out to answer two important research questions, the first being whether the PHY/MAC layers of LTE-V2X are sufficiently secure as to prevent exploitation by an intelligent attacker. We have definitively answered this question in the negative, showing through our two novel DoS attacks that significant improvements are needed to bring LTE-V2X up to this level of security. To this end, we proposed and validated detection techniques for each of our attacks—using an unsupervised cluster analysis algorithm and regression analysis, respectively—also; for each attack we laid out at least one promising mitigation approach as a basis for further study. These contributions addressed our second research question, which asked how we could effectively detect any attacks we identified despite the attacker's efforts to remain undetected.

In future work, we may investigate any of the various avenues we have indicated throughout this thesis. The specific method to be used by an attacker to subvert anonymization techniques and identify the sender of a particular BSM is well worthy of study and has broad applications beyond our scope of LTE-V2X DoS attacks. With respect to our targeted sidelink jamming attack, we intend to further analyze its effectiveness and determine whether, as anticipated, it is more stealthy in a realistic channel with considerations made for factors like multi-path fading and lower signal-to-noise ratio. Also, a more detailed study on how traffic patterns would affect the interval in which a monitor might be able to attempt detection of the attack would round out our discussion on that topic. For our sidelink resource exhaustion attack, we would be interested to evaluate how the additional considerations of upper-layer security protocols would help or harm the attacker's ability to remain undetected. Similarly, extending our attack design to factor in the resource-reservation capabilities that exist in some upper-layer V2V protocols may help increase the effectiveness of our attack. Finally, for both of our attacks, we intend to pursue the various mitigation techniques we have proposed with greater rigor, with the hope of ultimately devising a low-cost (in terms of system performance) mitigation for each attack to be applied to a revision of the LTE-V2X standard.

Appendix - Optimizing Sidelink Jamming through Abuse of Upper-Layer Security Protocols

This appendix describes the concept for an optimized version of the attack in Section 5 wherein Eve could transmit far less often and achieve the same results as our original attack. Since we were able to achieve experimental results that were promising, we include brief discussion of this optimization here along with a description of our obstacles and reasons for relegating it to this appendix rather than including it in the main thesis report. A bit of brief additional background is necessary to explain our approach. At the upper layers of V2V protocols, security is provided through the IEEE 1609.2 standards [38]. This standard provides, among other things, a requirement for the use of certificate-based message authentication using digital signatures. This is intended to mitigate upper-layer security threats like message spoofing, credential theft, message tampering, etc. Thus, each BSM is digitally signed using a certificate which must be validated by the receiving vehicle. In order for the receiver to validate the signature on a message, the receiver needs to obtain the signer's certificate in order to validate it, so BSMs include not only digital signatures but also the certificates used to generate them. However, certificates are relatively large and including them in every message adds latency, so vehicles only include their certificate in every fifth message [67]. The four messages between each certificate-bearing message include a SHA-256 digest of the certificate instead so that the receiver can look up the certificate received on a previous message to authenticate messages which do not include the certificate. Our idea, based on knowledge of these interwoven requirements, was to allow Eve to jam only those messages from Alice which include Alice's certificate. Since other vehicles would only receive the four out of every five messages from Alice which do not bear Alice's certificate, those vehicles would be unable to verify the messages' signatures and consequently would be forced to ignore them. This optimization, if successful, would improve Eve's efficiency significantly, allowing her to cause $>90\%$ packet loss for Alice while only actively jamming 20% of her messages.

We validated this approach using the same setup described in Section 5.2.1, confirming that this attack has the potential to be effective against commercial LTE-V2X equipment. However, we encountered a significant obstacle: a certificate-bearing message is often the first to be transmitted after Alice performs SPS resource reselection. Eve is not able to jam the first message after resource reselection (see Section 5), so this certificate would be distributed to all vehicles within range. Once Alice's certificate is distributed, Eve would need to fall back on the attack as described in Section 5.1 to remain effective. We analyzed the severity of this problem and found that it is a critical issue with this optimization: in over 85% of cases, Alice distributes her certificate on the first message after resource reselection within 15 seconds of Eve beginning her attack. As we were unable to overcome this obstacle, we present this appendix both due to the importance of negative results and because we believe this is not an entirely insurmountable obstacle. Some related work (e.g., [52]) has very recently looked at using machine learning techniques to accurately predict the resources a vehicle will reselect; if such an approach could be adopted, then this optimization may ultimately prove viable. For now, we leave such an investigation to future work.

References

[1] A. Boyaci, S. Yarkan, A. R. Ekti, and M. A. Aydin, “WWW: World Wide Wheels—A paradigm shift for transportation systems via xG,” in *Vehicle-to-Vehicle and Vehicle-to-Infrastructure Communications: A Technical Approach*, F. Hu, Ed. CRC Press, 2018, ch. 16, sec. 2.2.

[2] National Highway Traffic Safety Administration, “Technical report 11078-101414-v2a,” 2014, Accessed: Feb. 20, 2020. [Online]. Available: https://bit.ly/35EggyG

[3] M. Won, “A review on V2V communication for traffic jam management,” in *Vehicle-to-Vehicle and Vehicle-to-Infrastructure Communications: A Technical Approach*, F. Hu, Ed. CRC Press, 2018, ch. 1.

[4] A. Ali, L. Jiang, S. Patil, J. Li, and R. W. Heath, “Vehicle-to-vehicle communication for autonomous vehicles: Safety and maneuver planning,” in *Proc. IEEE Veh. Technol. Conf. (VTC)*, Chicago, IL, USA, Aug. 2018.

[5] S. Singh, “Connected car market worth $212.7 billion by 2027,” Nov. 2019, Accessed: May 27, 2020. [Online]. Available: https://www.marketsandmarkets.com/PressReleases/connected-cars.asp

[6] *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Wireless Access in Vehicular Environments*, IEEE Standard 802.11p, 2010.

[7] *Summary of Rel-14 Work Items*, 3rd Generation Partnership Project Technical Report 21.914 V14.0.0, 2018.

[8] *Summary of Rel-15 Work Items*, 3rd Generation Partnership Project Technical Report 21.915 V15.0.0, 2019.

[9] *Summary of Rel-16 Work Items*, 3rd Generation Partnership Project Technical Report 21.915 V16.0.0, 2020.

[10] Federal Communications Commission, “Dedicated short range communications (DSRC) service,” 2019, Accessed: February 21, 2020. [Online]. Available: https://www.fcc.gov/wireless/bureau-divisions/mobility-division/dedicated-short-range-communications-dsrc-service

[11] 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.

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

[13] S. Antipolis, “ETSI C-V2X plugtest achieves interoperability success rate of 94%,” Aug. 2020, Accessed: June 9, 2021. [Online]. Available: https://www.etsi.org/committee?id=1810

[14] 5G Automotive Association (5GAA), "A visionary roadmap for advanced driving use cases, connectivity technologies, and radio spectrum needs," Sep. 2020, Accessed: June 8, 2021. [Online]. Available: https://5gaa.org/news/the-new-c-v2x-roadmap-for-automotive-connectivity/

[15] D. Butler, "How 'talking' and 'listening' vehicles could make roads safer, cities better," Jan. 2019, Accessed: June 9, 2021. [Online]. Available: https://medium.com/cityoftomorrow/how-talking-and-listening-vehicles-could-make-roads-safer-cities-better-f215c68f376f

[16] Federal Communications Commission, "In the matter of use of the 5.850-5.925 GHz band (ET Docket No. 19-138)," Nov. 2020, Accessed: February 10, 2021. [Online]. Available: https://docs.fcc.gov/public/attachments/FCC-20-164A1.pdf

[17] Qualcomm Technologies, Inc., "Cellular V2X technology overview," Qualcomm Technologies, Inc.", 2019, Accessed: June 3, 2020. [Online]. Available: https://www.qualcomm.com/media/documents/files/c-v2x-technology-overview.pdf

[18] 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

[19] *Physical layer procedures*, 3rd Generation Partnership Project Technical Specification 36.213 V16.5.0, 2021.

[20] *Architecture enhancements for 5G System (5GS) to support Vehicle-to-Everything (V2X) services (Release 17)*, 3rd Generation Partnership Project Technical Specification 23.287 V17.0.0 (2021-06), 2021.

[21] 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.

[22] 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.

[23] 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.

[24] 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. Sensing, Commun. and Netw. (SECON)*, Virtual Conference, Jul. 2021.

[25] A. Lautenbach, N. Nowdehi, T. Olovsson, and R. Zaragatzky, "A preliminary security assessment of 5G V2X," in *Proc. IEEE Veh. Technol. Conf. (VTC)*, Kuala Lumpur, Malaysia, Apr. 2019.

[26] 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.

[27] V.-L. Nguyen, P.-C. Lin, and R.-H. Hwang, "Physical signal-driven fusion for V2X misbehavior detection," in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Los Angeles, CA, USA, Dec. 2019.

[28] 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.

[29] C. Wang, Z. Li, J. Shi, J. Si, and D. W. K. Ng, "Physical layer security of vehicular networks: A stochastic geometry approach," in *Proc. IEEE Int. Conf. Commun. Workshops (ICC Workshops)*, Dublin, Ireland, Jul. 2020.

[30] V. Marojevic, "C-V2X security requirements and procedures: Survey and research directions," Jul. 2018. [Online]. Available: http://arxiv.org/abs/1807.09338

[31] B. M. ElHalawany, A. A. A. El-Banna, and K. Wu, "Physical-layer security and privacy for vehicle-to-everything," *IEEE Commun. Mag.*, vol. 57, pp. 84–90, Oct. 2019.

[32] C. Wang, Z. Li, X.-G. Xia, J. Shi, J. Si, and Y. Zou, "Physical layer security enhancement using artificial noise in cellular vehicle-to-everything (C-V2X) networks," *IEEE Trans. Veh. Technol.*, vol. 69, pp. 15 253–15 268, 2020.

[33] Y. Liu, W. Wang, H.-H. Chen, L. Wang, N. Cheng, W. Meng, and X. Shen, "Secrecy rate maximization via radio resource allocation in cellular underlaying V2V communications," *IEEE Trans. Veh. Technol.*, vol. 69, pp. 7281–7294, Apr. 2020.

[34] 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.

[35] *Overview of 3GPP Release 12 V12.0.0 (2015-09)*, 3rd Generation Partnership Project Release 12 V12.0.0, 2015.

[36] *User Equipment (UE) radio transmission and reception*, 3rd Generation Partnership Project Technical Report 36.785 V14.0.0: Vehicle to Vehicle (V2V) services based on LTE sidelink, 2021.

[37] *V2X Communications Message Set Dictionary*, SAE Standard J2735E 2020-07, 2020.

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

[39] *Wireless Access in Vehicular Environments–Security Services for Applications and Management Messages - Amendment 1*, IEEE Standard 1609.2a-2017, 2017.

[40] *Wireless Access in Vehicular Environments–Security Services for Applications and Management Messages - Amendment 2*, IEEE Standard 1602.2b-2019, 2019.

[41] *Wireless Access in Vehicular Environments (WAVE)–Certificate Management Interfaces for End Entities*, IEEE Standard 1609.2.1-2020, 2020.

[42] G. Twardokus, J. Ponicki, S. Baker, P. Carenzo, H. Rahbari, and S. Mishra, "Targeted discreditation attack against trust management in connected vehicles," in *Proc. IEEE Int. Conf. Commun. (ICC)*, (Virtual Conference) Montreal, QC, Canada, Jun. 2021.

[43] National Security Agency, "COMINT and COMSEC: The tactics of 1914-1918, part I," *Cryptologic Spectrum*, vol. 2, no. 3, 1972.

[44] ——, "Tempest: A signal problem," *Cryptologic Spectrum*, vol. 2, no. 3, 1972.

[45] D. Paul, "The Navajo code talkers." Pittsburgh, PA, USA: Dorrance Publishing Company, 1998, ISBN-13: 978-0805945904.

[46] 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.

[47] Ò. Puñal, C. Pereira, A. Aguiar, and J. Gross, "Experimental characterization and modeling of RF jamming attacks on VANETs," *IEEE Trans. Veh. Technol.*, vol. 64, no. 2, pp. 524–540, May 2014.

[48] 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.

[49] S. Feng and S. Haykin, "Cognitive risk control for anti-jamming V2V communications in autonomous vehicle networks," *IEEE Trans. Veh. Technol.*, vol. 68, no. 10, pp. 9920–9934, 2019.

[50] P. Gu, C. Hua, R. Khatoun, and Y. Wu, "Cooperative anti-jamming relaying for control channel jamming in vehicular networks," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Singapore, Dec. 2017.

[51] 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.

[52] 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 Conference, Feb. 2021.

[53] N. Lyamin, D. Kleyko, Q. Delooz, and A. Vinel, "AI-Based malicious network traffic detection in VANETs," *IEEE Netw.*, vol. 32, no. 6, pp. 15–21, Nov. 2018.

[54] S. Hu, Q. A. Chen, J. Sun, Y. Feng, Z. M. Mao, and H. X. Liu, "Automated discovery of denial-of-service vulnerabilities in connected vehicle protocols," in *Proc. USENIX Security Symp.*, Virtual Conference, Aug. 2021.

[55] A. Yang, J. Weng, N. Cheng, J. Ni, X. Lin, and X. Shen, “DeQoS attack: Degrading quality of service in VANETs and its mitigation,” *IEEE Trans. Veh. Technol.*, vol. 68, no. 5, pp. 4834–4845, 2019.

[56] D. M. Mughal, J. S. Kim, H. Lee, and M. Y. Chung, “Performance analysis of V2V communications: A novel scheduling assignment and data transmission scheme,” *IEEE Trans. Veh. Technol.*, vol. 68, no. 7, pp. 7045 – 7056, 2019.

[57] 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.

[58] A. Abdelaziz, C. E. Koksal, R. Burton, F. Barickman, J. Martin, J. Weston, and K. Woodruff, “Beyond PKI: Enhanced authentication in vehicular networks via MIMO,” in *Proc. IEEE Int. Workshop Signal Process. Advances in Wireless Commun. (SPAWC)*, Kalamata, Greece, Jun. 2018.

[59] M. Sun, Y. Man, M. Li, and R. Gerdes, “SVM: Secure vehicle motion verification with a single wireless receiver,” in *Proc. ACM Conf. Secur. and Privacy in Wireless and Mobile Netw. (WiSec’20)*, Linz, Austria, Jul. 2020.

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

[61] T. Ebinuma, “GPS-SDR-SIM,” 2018, Accessed: April 20, 2020. [Online]. Available: https://github.com/osqzss/gps-sdr-sim

[62] 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. ACM Int. Workshop Wireless Netw. Testbeds, Experimental Eval., and Characterization (WINTECH)*, New York City, NY, USA, Oct. 2016, pp. 25–32.

[63] 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.

[64] Ettus Research, “B200/B210/B200mini/B205mini,” 2021, Accessed: July 19, 2021. [Online]. Available: https://kb.ettus.com/B200/B210/B200mini/B205mini#B210_2

[65] 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.

[66] J. A. Shaw, “Radiometry and the Friis transmission equation,” *American Journal of Physics*, vol. 81, no. 1, pp. 33–37, 2013.

[67] *Dedicated Short Range Communication (DSRC) Systems Engineering Process Guidance for SAE J2945/X Documents and Common Design Concepts*, SAE Standard J2945 2017-12, 2017.

[68] 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.

[69] New York State Department of Transportation, “Traffic data viewer,” 2021, Accessed: June 27, 2021. [Online]. Available: https://www.dot.ny.gov/tdv

[70] 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, Jul. 2017.

[71] H. K. Kanagala and V. J. R. Krishnaiah, “A comparative study of K-Means, DBSCAN and OPTICS,” in *Proc. Int. Conf. Comput. Commun. and Informatics (ICCCI)*, Coimbatore, India, Jan. 2016.

[72] 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, 2021.

[73] 5GAA, “Test procedures and results for 20-MHz deployment in CH183,” 2019, Accessed: July 21, 2021. [Online]. Available: https://bit.ly/2V8vF8z