Cross-Layer Performance Analysis of Post-Quantum Digital Signatures in Urban 5G-V2X Communications

Caio Teixeira, Marco Amaral Henriques

Departamento de Engenharia de Computação e Automação (DCA)
Faculdade de Engenharia Elétrica e de Computação (FEEC)
Universidade Estadual de Campinas (Unicamp) – Campinas, SP – Brazil

caio@dca.fee.unicamp.br maah@unicamp.br

Abstract. *As quantum computing threatens existing public key algorithms within the next decade, transitioning to post-quantum cryptography (PQC) is imperative. This paper evaluates the physical and MAC layer impacts of adopting PQC in 5G-V2X networks. Using the WiLabV2Xsim simulator in a 3GPP-compliant urban scenario, we assess network degradation across different post-quantum signature algorithms, and provide evidence that algorithms with large signatures cause severe radio channel saturation and latency, while compact alternatives introduce heavy computational bottlenecks. We conclude that direct ECDSA replacement is unfeasible, and that securing future V2X systems requires fundamental changes on the current PKI paradigm.*

1. Introduction

Post-quantum cryptography has been under standardization for almost a decade, with many countries establishing new alternatives for traditional public-key cryptography. However, a global transition from traditional to post-quantum cryptography is slow, and recent advances in both quantum computing and quantum attacks impose an even faster pace on the process. Furthermore, post-quantum algorithms bring a significant increase in communication cost, which can hinder widespread adoption [Alagic 2025].

In the US, the National Institute of Standards and Technology (NIST) has already begun enforcing this transition, deprecating traditional algorithms by 2030 and disallowing them by 2035 [Regenscheid 2024]. This timeline matches the planned deployment of 6G networks, which are currently under study by the mobile standards initiative 3GPP, thus requiring that these networks and their supported applications be quantum-safe from day one. However, the range of applications supported by mobile networks is wide, and they have significantly different requirements and limitations.

Within these applications supported by mobile networking are Vehicle-to-Everything (V2X) communications. Connected vehicles are becoming commonplace in many countries, and different visions for smart cities have already started being implemented across the world [Rito et al. 2023]. Future visions include autonomous driving in ground vehicles, alongside unmanned aerial vehicles (or drones) supporting logistics, communications and many other applications. These use cases demand a secure networking environment, as failures and attacks can have critical impact, causing great monetary loss or even lives.

Safety applications, such as those supporting both drivers and autonomous vehicles, require information shared through broadcast messages, and authenticating such

messages is critical to ensure no attackers can influence other vehicles. Currently, this authentication is done through digital signatures, using traditional methods. However, the transition to post-quantum algorithms severely affects these systems, as messages have short payloads, require strong authentication and, as smart cities grow, must support increasingly dense scenarios in a wireless radio channel.

Many works have assessed how current post-quantum standards might be integrated into current V2X technologies, either evaluating their performance on appropriate platforms or proposing strategies to mitigate the added cost through theoretical frameworks. However, little ground has been covered in assessing how impactful the choice of algorithm is in practice, especially when considering the characteristics of the physical layer, where congestion control and signal interference play an important role on how applications actually behave.

In this work, we focus on evaluating post-quantum signature algorithms for Cooperative Awareness Messages (CAMs) in an urban scenario, using 5G New Radio (5G-NR) as the radio technology. This combination, called NR-V2X (or 5G-V2X), is highly flexible and capable of operating in multiple modes, allowing a configurable trade-off between signal range and data rates. We evaluate how the choice of signature algorithm impacts the network by simulating communications through WiLabV2XSim [Todisco et al. 2021]. We use different channel configurations for each algorithm, tailoring the choice of Modulation and Coding Scheme (MCS) and sub-channel size to guarantee longer reception ranges while providing multiplexing capabilities to support dense scenarios in line with 3GPP evaluation methodologies. We identify the boundaries of service imposed by each algorithm, exposing the physical layer limitations brought by the adoption of post-quantum cryptography.

2. Fundamentals

2.1. Post-Quantum Cryptography and Standardization

Traditional cryptography has been under threat by quantum computers since 1994 due to Shor's algorithm [Shor 1994], which can break the underlying problems that guarantee their security. Quantum computing has been advancing fast, and new improvements on quantum attacks may accelerate this issue even further [Webster et al. 2026].

Post-quantum cryptography (PQC) encompasses the set of cryptographic algorithms that are resilient to such attacks by basing their security on problems believed to be hard for quantum computers to solve while also being secure in the classic paradigm. These algorithms can be categorized according to their subjacent mathematical problems. Most notably, lattice-based cryptography has been one of the most promising categories, as it presents high performance and is based on variants of well-studied problems [Alagic 2025]. Other categories have also seen recent improvements, with schemes based on multivariate equations and isogenies between elliptic curves achieving even lower communication costs [Alagic 2026].

While many countries have begun standardizing post-quantum algorithms, the US-based NIST has already selected for standardization two Key Encapsulation (KEM) and three Digital Signature algorithms. Of these, one KEM and two digital signature algorithms are lattice-based, while the remaining two are considered "fallbacks" in case

lattice problems are proven fragile in the future. Due to the fairly larger computation and communication costs of those fallbacks, NIST has opened an additional call for digital signature algorithms, which is now in its third round of evaluation [Alagic 2026].

While PQC solves the quantum attack problem, it poses new challenges as well. While computational costs vary between algorithms, the increase in public key and signature sizes brought by all of them have a severe impact from a communication standpoint. The main standard for traditional signatures, ECDSA, requires 64 bytes per signature, while FALCON, one of the algorithms under standardization and that has been designed to lower communication costs, achieves at best 666 bytes per signature [Fouque et al. 2020]. Some of the new alternatives from the additional call achieve smaller signatures and/or public key sizes, such as MAYO [Beullens et al. 2025] and SQISign [Aardal et al. 2025]; however, they require significantly more computing power for each operation, to the point where SQISign is unable to be executed in constrained environments except for signature verification.

Transitioning from traditional to post-quantum cryptography is not, then, simply a matter of replacing traditional algorithms. Each use case is impacted differently depending on the bottlenecks and constraints they face, so careful consideration must be taken in choosing which algorithm to adopt. In decentralized, time-critical environments such as vehicular networks, communication overheads can translate to either increased latency or complete saturation of the physical channel [Yoshizawa and Preneel 2023].

2.2. Vehicle-to-Everything (V2X) Communications

The term Vehicle-to-Everything (V2X) describes vehicular communications in a broader sense than just vehicle-to-vehicle, as connected vehicles can exchange data with infrastructure, pedestrians, over the network and so on. The standards that define the infrastructure and technology used for V2X vary from region to region, as such standards encompass not only technical decisions, but also regulatory ones. The most prominent standards today are DSRC (Dedicated Short-Range Communications) in the US and C-ITS (Cooperative Intelligent Transport Systems) in Europe [Sedar et al. 2023]. Both are based on the IEEE 802.11p standard for medium access and radio technologies, but for the network and up to the application layers, DSRC is defined by the IEEE 1609 family of standards (WAVE), while C-ITS is composed by a series of ETSI standards such as ETSI 103 301 and ETSI TS 102 941.

V2X encompasses a wide range of applications and use cases, such as road safety, traffic monitoring and efficiency management, platooning, smart city sensoring, remote driving, and many others. Road safety is one of the most critical of such applications. The core component of this application is Cooperative Awareness Messages (CAMs)[1]. CAMs are short messages each vehicle or infrastructure device broadcasts every 100ms to 1ms, containing information about the sender such as position, speed, acceleration, vehicle length and width, and other data used for safety applications. Another important type of safety messages is Decentralized Environmental Notification Messages (DENMs), which report events such as road blocks, weather hazards and other traffic events. In this work, we focus on CAMs due to their prevalence, frequency and time-sensitiveness.

[1]This is the name specified by ETSI; for DSRC, they are called Basic Safety Messages, or BSMs, but have similar structure and serve the same purpose.

CAMs contain information used for decision-making, especially if we consider future scenarios with autonomous vehicles that rely on sensing and cooperation; therefore, authentication is crucial. In current V2X standards, ECDSA has been adopted as the solution due to its relatively short size (64 bytes per signature), and every message is signed before being broadcast. The payload size of CAMs can vary from 41 bytes up to 400 bytes, depending on whether they contain optional/infrequent information (such as path history and special vehicle information) or not. Considering the overhead added by different layer headers and ECDSA signatures, a more frequent packet has around 190 bytes, while common less frequent packets are usually 300 bytes in size [ETSI 2019].

However, if we consider transitioning to quantum-safe alternatives, the impact of signatures on the total packet size becomes severe. Using FALCON (which is already under standardization, as discussed in Section 2.1), the cost of a signature increases tenfold, and the overhead introduced by authentication grows from an around 20% added packet size to over 200%. Even if network nodes have enough processing power to sign and verify such messages, the increase in size has direct impact on the shared physical channel used by V2X safety applications.

2.3. 5G-NR and NR-V2X

V2X communications may use different radio technologies for the medium access and physical layers. The IEEE 802.11p standard has been adopted by both DSRC and C-ITS; however the growing range of applications, especially those that require consistent network access, have shifted the interest towards mobile radio. This possibility was then concretized when 3GPP, the governing body for mobile network standards, released a specification for how mobile radio and network should support V2X specifications during Release 14 [ETSI 2025]. Since then, supported applications and use cases have evolved each release, each with their own performance indicators and supporting infrastructure.

This technology has been named Cellular V2X, or C-V2X, and can be further distinguished as either LTE-V2X, when using 4G technologies, or NR-V2X (sometimes called 5G-V2X), when using 5G New Radio (5G-NR). In this work, we will focus on NR-V2X as the MAC and PHY layer technology under analysis, as it lays the ground upon which future 6G technologies will be built.

For safety and cooperation, instead of relying on a central cell tower like mobile devices such as phones, vehicles instead use a second mode of communication - 5G Mode 2, or *Sidelink*. Sidelink enables communication between devices without intermediation from a cell tower, enabling services to run even on out-of-coverage areas and reducing broadcast latency. As CAMs are required to be broadcast to nearby vehicles and other actors, Sidelink becomes the ideal mode for this application.

At the physical layer, 5G-NR employs Orthogonal Frequency-Division Multiplexing, dividing the spectrum into subcarriers, which are then grouped in blocks of 12 subcarriers, called *Resource Blocks* (RBs) [Garcia et al. 2021]. NR-V2X communications happen over a 20 MHz channel, and depending on the *subcarrier spacing* (the “size” of each subcarrier), can support up to 106 RBs. Resource Blocks are further logically grouped into equally wide *subchannels*. Depending on the payload size, a vehicle must reserve one or multiple subchannels for a time slot to broadcast its message.

The amount of data that can be sent in a given time slot depends heavily on the

Modulation and Coding Scheme (MCS) chosen for the channel. Using more robust modulations and error-correcting codes means the transmission is more reliable and can travel further without being destroyed by interference; however, the amount of payload data sent during the time slot decreases. On the other hand, more aggressive modulations and lighter error-correcting codes leave more space for payload data, but are more susceptible to interference, therefore decreasing the range in which a receiver can successfully decode the message.

The primary 3GPP MCS table defines 29 configurations for 5G-NR, indexed from MCS 0 to MCS 28, ordered by increasing spectral efficiency. For NR-V2X, more aggressive modulations can be used to transmit bursts of payload data, such as when sharing LiDAR and video data; however, for safety applications, reliability is crucial, and more robust modulations are required [Garcia et al. 2021]. MCS 0 to 9 use QPSK modulation, and are therefore the standard for NR-V2X communications. Furthermore, to meet strict latency requirements, the physical layer is constrained to transmit the entire payload within a single time slot (typically 1ms). This time-slot limitation creates a scenario where signal robustness must be balanced against the multiplexing capabilities of the 20 MHz channel.

Packets signed with the traditional ECDSA algorithm fit into a fraction of the channel at low MCS. However, the increased size of post-quantum signatures forces a compromise: either adopt fragile, high-MCS configurations with shorter range, or adopt lower-MCS configurations that consume the majority of the channel's Resource Blocks, therefore resulting in higher probability of network congestion.

2.3.1. MAC Layer and Congestion Control

Since NR-V2X operates over Sidelink, there is no centralized cellular infrastructure to coordinate channel usage. Therefore, vehicles must autonomously allocate their own transmission resources using a decentralized algorithm known as *Sensing-Based Semi-Persistent Scheduling* (SB-SPS).

The SB-SPS algorithm operates in three phases. First, the vehicle continuously monitors the channel during a Sensing Window (usually the past second), decoding the control information of nearby vehicles and measuring interference. Then, when the vehicle needs to transmit a CAM, it projects a future Selection Window (e.g., the next 100ms) and excludes any Resource Blocks that have been already marked as reserved or exhibit high interference based on its sensing history. Finally, the vehicle randomly selects a block of subchannels from a pool of available resources. Since CAMs are generated periodically, the vehicle uses this selection to make a semi-persistent reservation, using the exact same subchannels for a set number of future broadcast intervals. Once this reservation counter expires, the vehicle repeats the sensing and selection process.

Due to the decentralized nature of SB-SPS, as the channel becomes congested, it becomes increasingly difficult for vehicles to select resource blocks without overlap. This forces the MAC scheduler to hold packets in transmission queues until resources become available. Furthermore, as the payload size increases, being able to allocate contiguous subchannels becomes highly unlikely. To prevent network collapse under such conditions,

Decentralized Congestion Control (DCC) is enforced [Garcia et al. 2021]. DCC defines a 60% Channel Busy Ratio (CBR) as a critical threshold, obligating vehicles to throttle their broadcast frequency or drop packets to alleviate channel saturation.

It is evident, then, that the increased communication cost of post-quantum signatures impacts the physical and MAC layers beyond the CPU or protocol bottlenecks traditionally evaluated in PQC literature. However, channel degradation is not uniform; it is highly dependent on different dynamic factors, including vehicle density (ρ), target transmission range, the presence of Line-of-Sight (LoS) versus Non-Line-of-Sight (NLoS) signal propagation and so on.

Since these parameters constantly change in real-world V2X environments, theoretical frameworks about the post-quantum impact cannot encompass its totality. It becomes necessary, then, to evaluate channel behavior under stress conditions to identify the exact boundaries where each post-quantum algorithm fails. By doing so, we shift from a binary definition of cryptographic adequacy to a more nuanced understanding of exactly which PQC algorithms can be used in vehicular networks and why.

3. Related Work

To the best of our knowledge, no other works in the literature have analyzed the impact of post-quantum algorithms in V2X down to the physical layer. However, there is literature on theoretical and protocol-level impacts of the post-quantum transition, as well as other works that evaluate physical layer configurations of NR-V2X for different use cases.

[Twardokus et al. 2024] evaluates post-quantum signatures in the DSRC framework, using with 802.11p technology. The authors considers the 2304-byte frame size of 802.11p as the packet size ceiling, and note that, even using FALCON, messages containing both signature and certificates exceed that limit. The article then proposes a partially-hybrid construction, where messages are signed with short-lived ECDSA keys, which are in turn authenticated by long-term, hybrid certificates with FALCON and ECDSA. The authors also note that adapting C-V2X to PQC would require significant changes beyond just swapping signature algorithms.

[Lonc et al. 2023] provides a thorough benchmark of post-quantum algorithms in the context of V2X applications, evaluating their execution times across different platforms, each representing a device profile of V2X networks. The analysis is done through the lens of ITS-G5 technology adopted by C-ITS, which, at its current state, imposes limitations on packet size right at the network layer that are stricter than its American counterpart DSRC. The authors note that FALCON and Kyber are the only algorithms that might fit the packet size constraints, especially through hybrid constructions; however, they note that balancing the trade-off of bandwidth and performance in this scenario will imply in modifications of current standards.

In the topic of NR-V2X configuration analysis, [Yan and Härri 2022] analyses different MCS configurations for NR-V2X Sidelink, considering a variety packet sizes, transmit rates and density. The authors find that MCS 24 outperforms other configurations in most cases, and that other MCS configurations can be explored depending on the application. However, these results consider a scenario with no mobility, consistent LoS and short range, which does not represent the majority of real-world V2X scenarios, therefore being applicable to only a small portion of V2X communications.

Finally, [Hegde et al. 2023] investigates the impact of the variable message sizes of Collective Perception Messages (CPM) on radio resource allocation for NR-V2X, using the same bandwidth allocated for CAMs. The authors find that the unpredictability and high ceiling of CPM packet sizes (50 to 800 bytes) drastically impact resource allocation, especially in dense scenarios, leading to outdated reception and high collision rates. These findings also give a baseline for the behavior of CAMs with large signatures, even though the variability is drastically reduced in an awareness and safety context.

4. Evaluation Methodology

To evaluate the physical-layer impact of post-quantum cryptographic overhead, we conducted system-level simulations of an urban use case of NR-V2X using the open-source discrete-event simulator WiLabV2Xsim [Todisco et al. 2021], across different vehicle density configurations. It has been employed in hundreds of previous works to benchmark both LTE-V2X and NR-V2X, as well as channel coexistence between mobile radio and 802.11p radio technologies [Wu et al. 2023]. Instead of relying on abstract network models, WiLabV2Xsim emulates the physical communication down to subframes at every simulation step, and features extensive parameterization across PHY and MAC layers, such as subchannel dimensions, MCS and DCC mechanisms, and many more.

For our purpose, the subframe-level emulation enables tracking of Resource Block (RB) allocation, channel sensing, and mid-air packet collisions, making it ideal for evaluating the capacity bottlenecks caused by large payloads. Furthermore, to guarantee that the results reflect real-world constraints, it models radio signal propagation according to the 3GPP guidelines for V2X evaluation [3GPP 2019], which provide mathematical models for path loss, shadowing, and fast-fading, allowing the simulator to dynamically calculate the signal degradation between vehicles as they move. It also enables simulation of predefined 3GPP-compliant scenarios, such as highway and urban communications, and accepts custom scenario descriptions defined through trace files and mobility models.

The urban scenario represents a critical point of evaluation for our investigation, as vehicle density may be very high (e.g. in a traffic jam), and V2X systems must be able to operate even under these conditions. An important aspect for this scenario is that the simulator constantly evaluates whether the communication link between two vehicles is Line-of-Sight (LoS) or Non-Line-of-Sight (NLoS). While LoS conditions represent a clear, unobstructed signal path (e.g. vehicles driving on the same street), NLoS conditions occur when the radio wave is physically obstructed by infrastructure, such as concrete buildings at an intersection, suffering from signal attenuation and multipath fading.

To evaluate the impact of PQC overhead in dense, realistic environments, we adopted the urban Manhattan grid scenario defined by the 3GPP guidelines [3GPP 2019]. This scenario simulates a 3x3 Manhattan grid with rectangular blocks, where vehicles on the same street have LoS, but have NLoS links with vehicles in other streets on intersections due to concrete buildings blocking the path. Figure 1 shows a visual representation of the scenario.

Our experiments were conducted under four different density configurations $\rho = \{25, 50, 100, 140\}$, where ρ represents the density of 5m-long vehicles per lane per km [3GPP 2019]. The total grid length, considering every horizontal and vertical street, is

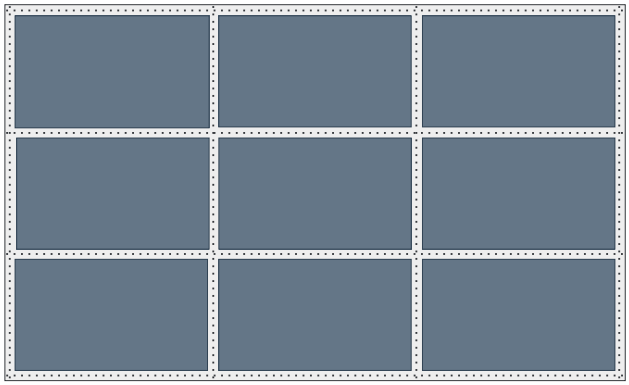

Figure 1. Representation of the simulated urban scenario, as specified in the 3GPP guidelines.

approximately 8.2km long. Since each street has 2 lanes, the sparse scenario with $\rho = 25$ has a total of $25 \times 8.2 \times 2 \approx 410$ vehicles on the grid, with an average of 35 meters of spacing between vehicles, simulating regular urban traffic. On the other hand, the most dense scenario has a total of $140 \times 8.2 \times 2 = 2296$ vehicles on the grid, with an average of 2.14m of spacing between them, accurately modeling an urban traffic jam scenario.

We assume CAMs are transmitted once every 100ms (10Hz) and the minimum possible payload (41 bytes total) for each message, representing a best-case scenario for providing collision warning, platooning and other services [3GPP 2019]. This allows us to establish a baseline for further evaluations and filter out options that cannot achieve even the lowest requirements. We also assume a total of 88 bytes of overhead due to L2 to L4 and SPDU headers, totaling 129 bytes for each packet, excluding the signature itself. We also do not consider certificates under this analysis, as they can be distributed under the standardized Peer-to-Peer Certificate Distribution method [ETSI 2021]; therefore, in an optimistic scenario, considering only signatures inside the packets is enough.

In Table 1, we present the simulated channel configurations, as well as total packet size, for each chosen algorithm, alongside a baseline ECDSA configuration for comparison. Channel configurations were defined heuristically to balance two competing physical constraints: maximizing subchannel multiplexing (to support high density) while strictly maintaining robust QPSK modulation (to ensure NLoS survivability). To achieve this, we have iterated on every possible combination of subchannel size and MCS, evaluating the amount of resource blocks required to send a signed packet for a given algorithm and how many vehicles would be able to use the channel simultaneously. Whenever different MCS configurations would yield the same multiplexing capacity, we chose the lowest MCS index to prioritize signal robustness, respecting the limitation of using only QPSK modulation, as described in Section 2.3.

We evaluate FALCON, under standardization as FN-DSA, as well as new algorithms of the additional NIST call that can achieve smaller signatures. We chose to evaluate MAYO, a multivariate equation-based algorithm in two configurations - MAYO_1, which balances public key and signature size, and MAYO_2, which has a larger public key, but minimizes signature size. Both offer a Security Level I, as defined by NIST. Finally, we also evaluate SQISign, an isogeny-based algorithm that achieves the shortest signature

Table 1. Channel configurations for each of the selected algorithms. *RBs per TB* denotes the amount of resource blocks (RBs) required to emit a single transport block (TB), which contains the CAM. *Multiplexing* denotes the number of parallel CAM transmissions supported by the configuration.

Algorithm	Packet Size	MCS	RBs per TB	Subchannel Size	Multiplexing
ECDSA	193 B	8	12	12 RBs	8
FALCON	795 B	7	50	10 RBs	2
$MAYO_1$	583 B	9	30	10 RBs	3
$MAYO_2$	315 B	8	20	10 RBs	5
SQISign	277 B	9	15	15 RBs	7

of the algorithms of the additional NIST call [Alagic 2026].

Notably, the primary NIST standard, ML-DSA, was excluded from our simulation. With a signature size of 2420 bytes, ML-DSA exceeds the maximum Transport Block Size (TBS) possible under QPSK modulations in a standard 20 MHz channel. Accommodating ML-DSA requires more aggressive modulations (i.e. 16-QAM or 64-QAM) which strip the signal of essential error correction, drastically increasing the likelihood of failure in NLoS urban environments. Therefore, it is physically unsuitable for safety broadcasts.

To evaluate our results, we look at different Key Performance Indicators (KPIs) for each simulation. We evaluate Packet Reception Ratio (PRR) from two perspectives. First, we analyze PRR as a function of distance for a set vehicle density, to understand how the choice of algorithm impacts transmission range. Then, we analyze PRR at a critical fixed reception range of 50 meters. As defined by 3GPP TR 22.885 [3GPP 2015], for urban use cases such as Forward Collision Warning, the 50-meter threshold mathematically corresponds to a mandatory 4-second driver response window at standard urban speeds of 45 km/h. By isolating this threshold, we evaluate whether each algorithm maintains the resiliency required to prevent collisions under fluctuating network densities.

We also analyze the Channel Busy Ratio (CBR) to quantify MAC-layer saturation. A high CBR indicates severe resource contention, triggering Decentralized Congestion Control (DCC) mechanisms that force vehicles to drop packets or throttle broadcast frequencies (as discussed in Section 2.3.1), directly compromising safety applications.

Finally, we evaluate latency through two distinct metrics: Update Delay and Data Age. Update Delay measures the transit time of a message from the MAC layer of the sender to a successful reception, isolating the time spent waiting in MAC-layer queues. Conversely, Data Age measures the continuous time elapsed since the receiver's last successful update from a specific sender. While Update Delay tracks network speed and whether vehicles can find available channels to send their messages in time, Data Age acts as a proxy for application-level safety, capturing the compounding penalty of packets lost to mid-air collisions.

5. Simulation Results

5.1. Packet Reception Ratio by Inter-Vehicle Distance

First, in Figure 2, we show the Packet Reception Ratio (PRR) by vehicle distance across the four selected densities to illustrate the impact of each configuration on the signal

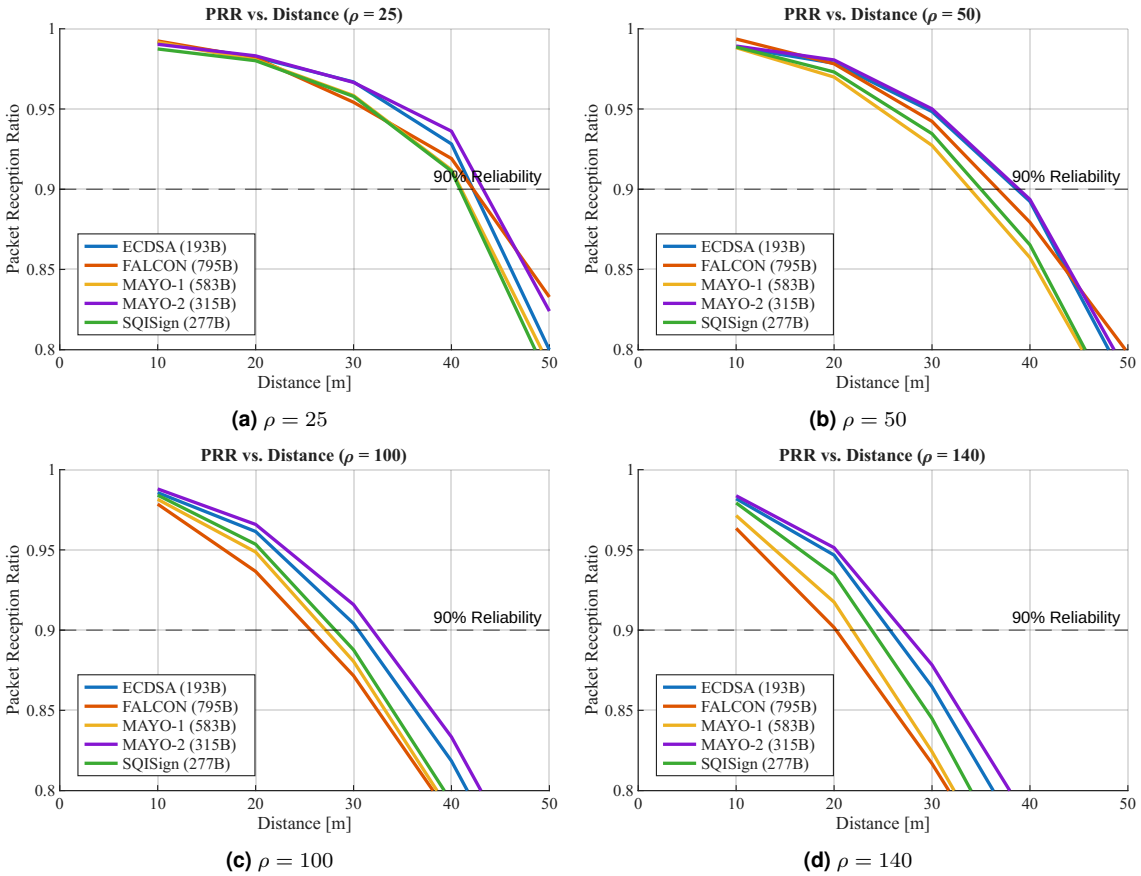

(a) $\rho = 25$ **(b)** $\rho = 50$

(c) $\rho = 100$ **(d)** $\rho = 140$

Figure 2. Packet Reception Ratio by inter-vehicle distance, across four different density configurations.

reception. Since FALCON uses a more robust MCS (MCS 7), it outperforms the other algorithms on sparse scenarios, degrading slower as the signal travels. The minimum reliability specified by 3GPP for collision warning [3GPP 2015] is 90% at effective distance; therefore, in the lowest density scenario, every algorithm is suitable until 40m at $\rho = 25$, but only MAYO_2 and ECDSA continue meeting this threshold at $\rho = 50$.

At higher densities, the increase in packet collisions due to hidden node transmissions sharply reduces PRR across every algorithm. However, this problem affects larger packets (e.g. FALCON and MAYO_1) even more, since they require more subchannels to transmit data, and the likelihood of a hidden node picking at least one of the same subchannels increases drastically. On the other hand, smaller packets that provide higher multiplexing (e.g. ECDSA and MAYO_2) are transmitted over fewer subchannels, therefore mitigating the chance of collisions due to hidden nodes.

5.2. Channel Busy Ratio

Figure 3 shows the increase in Channel Busy Ratio (CBR) as density increases. As mentioned in Section 2.3.1, due to the decentralized nature of SB-SPS and DCC, a 60% CBR represents a highly congested channel where packets suffer high collision rates and scheduling delays. This congestion threshold is highlighted in the figure.

The current standard ECDSA reaches slightly above the 60% CBR mark at the

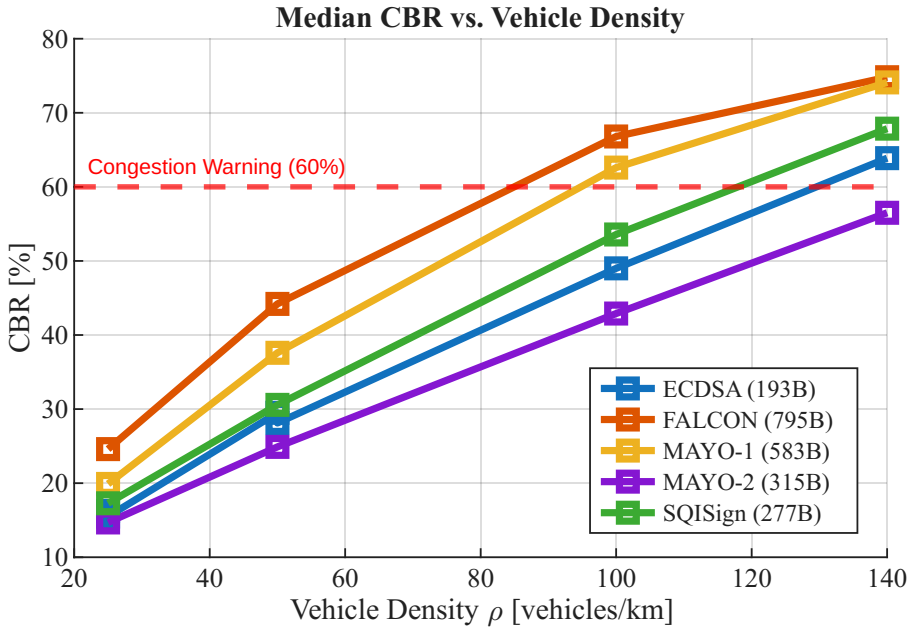

Figure 3. Median Channel Busy Ratio (CBR) by vehicle density.

highest density (63.8% CBR), but remains under reasonable expectations and does not collapse the network. For the more costly candidates, the impact of low multiplexing capacities becomes evident. FALCON reaches the 60% CBR threshold even below $\rho = 100$, and both it and MAYO_1 reach a 74% threshold at the highest density ($\rho = 140$). This behavior evidences a breakdown in communication, as many packets are being dropped at the MAC layer due to vehicles being unable to find empty subchannels for transmission.

An initially counterintuitive finding is that MAYO_2 consistently registers a lower median CBR than the ECDSA baseline, despite its larger payload size and identical MCS configuration. This inversion is a consequence of Power Spectral Density (PSD) dilution. In NR-V2X, total transmit power is distributed evenly across all allocated Resource Blocks. While the ECDSA payload concentrates its transmit power into a single 12-RB subchannel, MAYO_2 must span its payload across two 10-RB subchannels, consequently halving the transmission power per subchannel. This diluted signal falls below the sensing threshold at a much shorter physical distance, meaning fewer neighboring vehicles perceive the subchannels as occupied. While this artificially suppresses the network's perceived CBR, it comes at the cost of signal fragility and increased MAC-layer deferrals, as the scheduler struggles to allocate larger contiguous resource blocks in dense traffic.

Finally, SQISign performs slightly worse than ECDSA, but does not suffer from PSD dilution as MAYO_2 since every packet requires 15 RBs, and our subchannel configuration for this algorithm allocates exactly 15-RB wide subchannels.

5.3. Update Delay and Data Age

Figure 4 shows the variation of Update Delay and Data Age for two set inter-vehicle distances - namely, 50m and 100m. As mentioned before, Update Delay measures the time between the message being ready for transmission at the MAC layer and successful reception. Since the resource allocation window in NR-V2X is 100ms, the ideal scenario

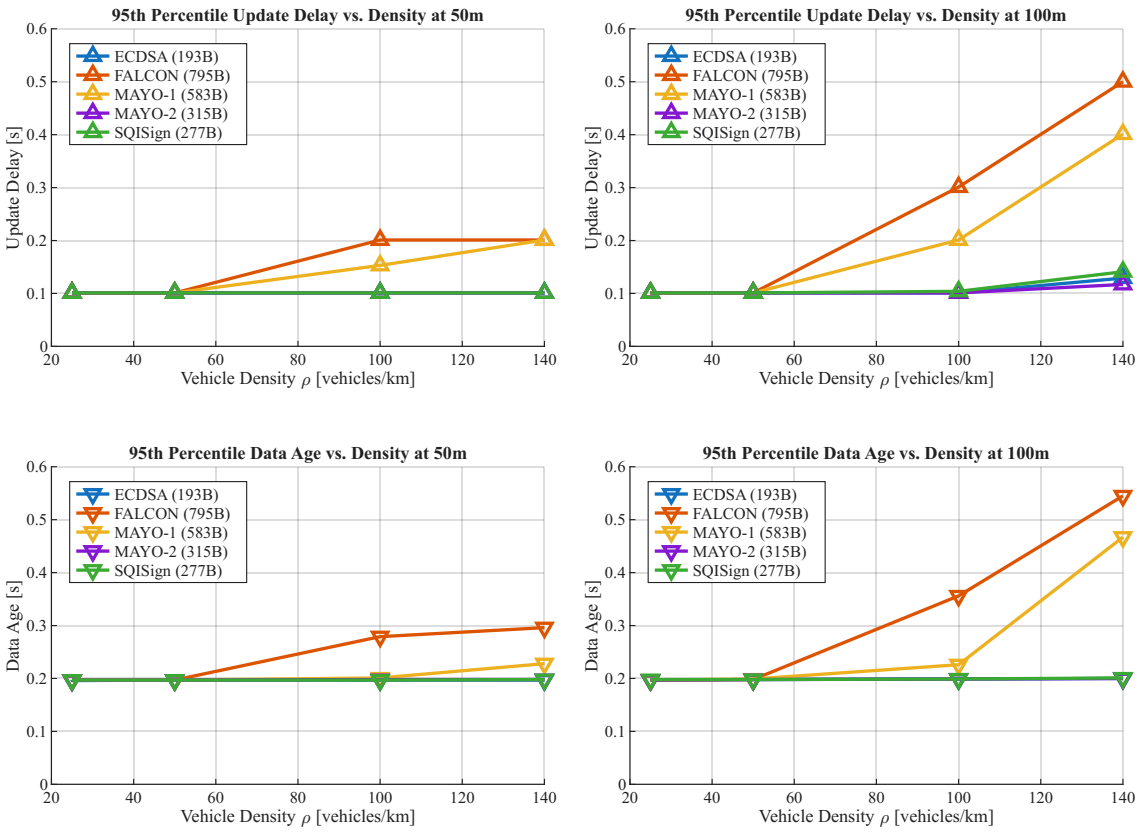

Figure 4. Update Delay and Data Age by vehicle density, sampled at either 50m or 100m of inter-vehicle distance.

would be for messages to have a consistent 100ms Update Delay.

However, at higher densities, the larger packets - namely, FALCON and MAYO_1 - start having difficulties allocating resources through SB-SPS, which leads to packets being retained at the MAC-layer queue. The critical problem, however, is that as messages are delayed, their payload information gets stale, which in turns means the whole decentralized awareness system starts getting outdated. This outdated information, then, either feeds decision-making algorithms, providing erroneous data, or is dropped due to lateness, starving decision-making algorithms from important information, especially from vehicles outside of line-of-sight, LiDAR range or other sensing mechanisms.

Data Age, on the other hand, represents how stale information about a certain vehicle can get in the network. If a vehicle cannot send messages, or has high packet loss due to collision (which is unrecoverable in a broadcast link due to lack of retransmission), other vehicles become effectively blind to its updates, also severely impacting moment-to-moment decision-making.

While 50m is the critical threshold for collision warning and avoidance, 100m is the standard range for cooperative awareness [ETSI 2019]. Therefore, the performance of both MAYO_1 and FALCON at 100m, achieving up to 400ms and 500ms of Update Delay along with 467ms and 545ms Data Age, respectively, is far beyond acceptable for V2X applications. Furthermore, we can notice that the remaining algorithms also suffer a

small penalty at high densities due to high channel congestion; however, they still remain at reasonable performance under both metrics.

6. Discussion

Given our simulation results and analysis in the previous section, we can affirm that FALCON, the only algorithm under standardization evaluated in this work, is unsuitable for V2X communications. While under low stress the algorithm performs well enough, it quickly degrades as density increases, causing application breakdowns well before maximum density. This poses a long-term risk in deploying FALCON as the post-quantum alternative for CAM signatures, as more connected vehicles grow more and more common and smart city systems are implemented around the world.

Regarding the additional signature schemes, both $MAYO_2$ and SQISign provide interesting alternatives to ECDSA in this context. In particular, the configuration chosen for $MAYO_2$ threads a fine line between multiplexing and channel congestion. While sending a $MAYO_2$-signed CAM requires two contiguous available subchannels, vehicles are still able to allocate resources even with high traffic. On the other hand, while PSD dilution may mask the CBR measurements, meaning the channel might be busier than the median leads us to believe, it does not affect either Update Delay or Data Age, meaning the configuration sits at the limit of the system's capacity without crashing it.

However, to achieve the goal of small signatures, both MAYO and SQISign have the trade-off of being very computationally costly. V2X cryptographic operations are usually done inside Hardware Security Modules, which may have constrained processing power compared to general-purpose CPUs. As an example, the ARM SecurCore SC300, a processor commonly deployed in vehicular HSMs, is built upon ARM Cortex-M3 cores. Furthermore, as smart cities expand V2X capabilities to more types of devices - especially aerial vehicles, such as drones - a high processing cost can become a major limitation.

In fact, at the time of writing, SQISign does not support key generation or signing on constrained devices (such as ARM Cortex-M4) [Aardal et al. 2025], and even just verifying takes up to seconds in processing time, along with a large memory footprint. The same goes for MAYO - while the algorithm supports constrained platforms, the authors evaluate that a single execution of $MAYO_2$ signature generation on ARM Cortex-M4 takes up to 100ms by itself, and around 50ms for verification [Beullens et al. 2024]. This computation time means the processor must spend a complete CAM emission cycle just signing the message from the previous cycle, and also means that no more than two CAMs would be able to be verified the same time frame - besides being mutually exclusive if the same core is used for both operations. While potential improvements to MAYO's verification times have been proposed in the literature [Boschini et al.], it is yet unclear how beneficial the technique is and whether such changes might be adopted in future iterations of the algorithm during the standardization process.

We conclude, then, that none of the evaluated options are able to meet the requirements of V2X networks without fundamental changes to its security paradigm. This work considers only a direct replacement of ECDSA by a post-quantum algorithm while keeping the same Public Key Infrastructure (PKI) with no further alterations. However, other authentication methods that replace the current PKI-based one can drastically change the scenario, as long as they provide alternatives to signing every

message with the algorithms evaluated in this work. One such method could be schemes based on TESLA and its variants, which rely on public key cryptography only to establish a commitment, while the packets themselves are authenticated through symmetric cryptography (such as HMACs) and timing-based constraints. While techniques such as pairing TESLA with digital signatures have already been evaluated as beneficial in multiple scenarios, including with Verify-on-Demand schemes [Cominetti et al. 2023], such techniques would be similarly affected by the signature size increase of PQC. However, if such algorithms are considered under the quantum threat scenario and a solution to distribute the signed secure anchor required by the algorithm is developed, they could prove to be suitable alternatives for securing V2X communications in the future.

7. Conclusion and Future Work

In this work, we have evaluated the impact of the transition to post-quantum signatures in the current V2X communication scenario by measuring the behavior of the physical layer using the WiLabV2XSim simulator [Todisco et al. 2021]. Through extensive analysis under different KPIs, we have shown that the impact on the physical layer is higher than a simple theoretical, protocol-level analysis could encompass. Given our discussion in Section 6, we provide the following takeaways:

1. **Current standards are nonviable as a simple replacement:** Through our analysis of FALCON's low performance for authenticating CAMs, along with the unsuitability of ML-DSA due to its large signature size, we conclude that current post-quantum signature standards do not meet the requirements for cooperative awareness and safety in V2X if they simply replace the current algorithms.
2. **New candidates may be viable only if heavily improved upon:** From our choice of candidates from the additional NIST call, we conclude that some of the post-quantum signature algorithms from the additional NIST call, e.g. MAYO, can be used in V2X scenarios as long as future optimizations and hardware implementations enable such algorithms to run efficiently on constrained devices.
3. **Structural changes might be required:** Finally, we conclude that further research on alternatives to the current PKI is required before transition to PQC becomes mandatory, along with research on alternative broadcast authentication schemes that mitigate the costs brought by PQC adoption.

As future work, different scenarios can be evaluated using the same methodology adopted in this paper, such as highways or scenarios with connected infrastructure nodes, as well as mixed CAM sizes, more traffic densities, and certificates management. While we have focused on high density scenarios, the high mobility of highways may yield different results that must also be considered when deciding how to handle the transition. Furthermore,real-world experiments with HSMs would provide further data on the channel behavior and how impactful the hardware limitations mentioned in Section 6. Finally, expanding the simulator to enable evaluation of alternatives to PKI, such as TESLA-based algorithms, can provide fruitful insights on the adequacy of these solutions to V2X communications.

Acknowledgements

This work was supported by Ericsson Telecomunicações Ltda., and by the Sao Paulo Research Foundation (FAPESP), grant 2021/00199-8, CPE SMARTNESS. The Gemini 3.1 Pro AI was used for proofreading and support in code debugging.

References

3GPP (2015). Study on LTE support for Vehicle to Everything (V2X) services. Technical Report (TR) 22.885, 3rd Generation Partnership Project (3GPP). Version 14.0.0.

3GPP (2019). Study on evaluation methodology of new Vehicle-to-Everything (V2X) use cases for LTE and NR. Technical Report (TR) 37.885, 3rd Generation Partnership Project (3GPP). Version 15.3.0.

Aardal, M. A., Adj, G., Aranha, D. F., Basso, A., Canales Martínez, I. A., Chávez-Saab, J., Corte-Real Santos, M., Dartois, P., De Feo, L., Duparc, M., Eriksen, J. K., Fouotsa, T. B., Gazzoni Filho, D. L., Hess, B., Kohel, D., Leroux, A., Longa, P., Maino, L., Meyer, M., Nakagawa, K., Onuki, H., Panny, L., Patranabis, S., Petit, C., Pope, G., Reijnders, K., Robert, D., Rodríguez-Henríquez, F., Schaeffler, S., and Wesolowski, B. (2025). SQIsign. Technical report, National Institute of Standards and Technology.

Alagic, G. (2025). Status Report on the Fourth Round of the NIST Post-Quantum Cryptography Standardization Process. Internal Report (IR) IR 8545, NIST.

Alagic, G. (2026). Status Report on the Second Round of the Additional Digital Signature Schemes for the NIST Post-Quantum Cryptography Standardization Process. Internal Report (IR) IR 8610, NIST.

Beullens, W., Campos, F., Celi, S., Hess, B., and Kannwischer, M. J. (2024). Nibbling mayo: Optimized implementations for avx2 and cortex-m4. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2024(2):252–275.

Beullens, W., Campos, F., Celi, S., Hess, B., and Kannwischer, M. J. (2025). MAYO specification document. Submission to the NIST Post-Quantum Cryptography Additional Digital Signatures Call. Available at `https://pqmayo.org/`.

Boschini, C., Fiore, D., Pagnin, E., Torresetti, L., and Visconti, A. Progressive and efficient verification for digital signatures: Extensions and experimental results. 14(3):551–575.

Cominetti, E. L., Silva, M. V. M., Simplicio, M. A., Kupwade Patil, H., and Ricardini, J. E. (2023). Faster verification of V2X basic safety messages via Message Chaining. *Vehicular Communications*, 44:100662.

ETSI (2019). Specification of cooperative awareness basic service. European Standard ETSI EN 302 637-2, ETSI.

ETSI (2021). Security header and certificate formats. Technical Specification ETSI TS 103 097, ETSI.

ETSI (2025). Service requirements for v2x services. Technical Specification ETSI TS 122 185, ETSI.

Fouque, P.-A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Prest, T., Ricosset, T., Seiler, G., Whyte, W., and Zhang, Z. (2020). Falcon: Fast-Fourier lattice-based compact signatures over NTRU. Submission to the NIST Post-Quantum Cryptography Standardization Project.

Garcia, M. H. C., Molina-Galan, A., Boban, M., Gozalvez, J., Coll-Perales, B., Şahin, T., and Kousaridas, A. (2021). A Tutorial on 5G NR V2X Communications. *IEEE Communications Surveys & Tutorials*, 23(3):1972–2026.

Hegde, A., Delooz, Q., Mariyaklla, C. L., Festag, A., and Klingler, F. (2023). Radio Resource Allocation for Collective Perception in 5G-NR Vehicle-to-X Communication Systems. In *2023 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 1–7. ISSN: 1558-2612.

Lonc, B., Aubry, A., Bakhti, H., Christofi, M., and Mehrez, H. A. (2023). Feasibility and Benchmarking of Post-Quantum Cryptography in the Cooperative ITS Ecosystem. In *2023 IEEE Vehicular Networking Conference (VNC)*, pages 215–222. ISSN: 2157-9865.

Regenscheid, A. (2024). Transition to Post-Quantum Cryptography Standards. Technical Report NIST IR 8547 ipd, National Institute of Standards and Technology, Gaithersburg, MD.

Rito, P., Almeida, A., Figueiredo, A., Gomes, C., Teixeira, P., Rosmaninho, R., Lopes, R., Dias, D., Vítor, G., Perna, G., Silva, M., Senna, C., Raposo, D., Luís, M., Sargento, S., Oliveira, A., and de Carvalho, N. B. (2023). Aveiro Tech City Living Lab: A Communication, Sensing, and Computing Platform for City Environments. *IEEE Internet of Things Journal*, 10(15):13489–13510.

Sedar, R., Kalalas, C., Vázquez-Gallego, F., Alonso, L., and Alonso-Zarate, J. (2023). A Comprehensive Survey of V2X Cybersecurity Mechanisms and Future Research Paths. *IEEE Open Journal of the Communications Society*, 4:325–391.

Shor, P. (1994). Algorithms for quantum computation: discrete logarithms and factoring. In *Proc. 35th Annual Symposium on Foundations of Computer Science*, pages 124–134.

Todisco, V., Bartoletti, S., Campolo, C., Molinaro, A., Berthet, A. O., and Bazzi, A. (2021). Performance Analysis of Sidelink 5G-V2X Mode 2 Through an Open-Source Simulator. *IEEE Access*, 9:145648–145661.

Twardokus, G., Bindel, N., Rahbari, H., and McCarthy, S. (2024). When Cryptography Needs a Hand: Practical Post-Quantum Authentication for V2V Communications. In *Proceedings 2024 Network and Distributed System Security Symposium*, San Diego, CA, USA. Internet Society.

Webster, P., Berent, L., Chandra, O., Hockings, E. T., Baspin, N., Thomsen, F., Smith, S. C., and Cohen, L. Z. (2026). The Pinnacle Architecture: Reducing the cost of breaking RSA-2048 to 100 000 physical qubits using quantum LDPC codes. arXiv:2602.11457 [quant-ph].

Wu, Z., Bartoletti, S., Martinez, V., Todisco, V., and Bazzi, A. (2023). Analysis of Co-Channel Coexistence Mitigation Methods Applied to IEEE 802.11p and 5G NR-V2X Sidelink. *Sensors*, 23(9).

Yan, J. and Härri, J. (2022). MCS Analysis for 5G-NR V2X Sidelink Broadcast Communication. In *2022 IEEE Intelligent Vehicles Symposium (IV)*, pages 1347–1352.

Yoshizawa, T. and Preneel, B. (2023). Post-Quantum Impacts on V2X Certificates – Already at The End of The Road. In *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, pages 1–6.